

THE No 1 UK MAGAZINE FOR ELECTRONICS TECHNOLOGY & COMPUTER PROJECTS


EPE EVERYDAY PRACTICAL ELECTRONICS

www.epemag.co.uk

www.epemag.com

LEAD-ACID BATTERY ZAPPER

**Get extra life from your tired
old lead-acid battery**



MINICAL 5V METER CALIBRATION STANDARD **Plus a frequency reference/crystal checker**




**VIDEO
READING AID**
**Enlarges text and
displays it on a TV
or monitor**

- **USING MPLAB – 2**
- **MECHATRONICS – Controlling Motors**

\$6.95 US \$8.99 CAN
JULY 2007 PRINTED IN THE UK





By JIM ROWE

Lead-Acid Battery Zapper!

This simple circuit is designed to extend the working life of liquid-electrolyte lead-acid batteries, by dissolving the lead-sulphate crystals which form on their plates. It's powered by the battery itself (or by a charger) and 'zaps' the battery with a series of high-voltage pulses.

LEAD-ACID BATTERIES have been around for over 170 years now – ever since Gaston Plante built the first one back in 1834. They are used in huge numbers all around the world, mainly in the automotive industry. There's at least one in virtually every car, truck and bus to start the engine and power ancillary equipment. Multiple lead-acid batteries are also used in many electric vehicles to provide the motive power.

They're also used in large numbers for energy storage in solar and wind power plants. And by the way, we're talking about 'wet' or liquid electrolyte batteries here (also called 'flooded' lead-acid batteries).

The lead-sulphate effect


Although we'd now be lost without them, lead-acid batteries are not without their faults. Probably their main

drawback is that they have a relatively short working life, typically no more than about three or four years.

Why is this? Well, every time energy is drawn from a lead-acid battery, lead and sulphate ions from the electrolyte combine and are deposited on the plates in the form of soft lead-sulphate crystals. Then, when the battery is recharged, these crystals dissolve again in the sulphuric acid electrolyte.


More accurately, **MOST** of them re-dissolve – but not all. Even if the battery is never over-discharged and is always recharged promptly after it has been discharged, a small proportion of the lead sulphate remains on the plates. These then harden into 'hard' lead-sulphate crystals which are much less soluble and less conductive than before.

In practice, the formation of these hard lead-sulphate crystals




(A) CAPACITOR CHARGING PHASE

Fig.1(a): during the first phase of the circuit's operation, current flows from the battery (or charger) and charges a 100μF electrolytic capacitor via inductor L2.



(B) ENERGY TRANSFER PHASE

Fig.1(b): next, the switch is closed for 50μs, and current flows from the capacitor into L1. As a result, the energy stored in the capacitor is transferred to the inductor's magnetic field.



(C) DISCHARGE/PULSE PHASE

Fig.1(c): finally, the switch opens again, interrupting the inductor current and causing a high-voltage pulse across the inductor with the polarity shown. The green arrow shows the discharge current path.

gradually reduces the energy storage capacity of the battery. It does this both by masking the active areas on the plates and also by reducing the concentration of lead and sulphate ions in the electrolyte.

This 'sulphation' effect has been understood for many years. It's also well known that the effect occurs much faster if a battery is over-discharged, left in a discharged state for more than a few hours, or frequently under charged. In fact, batteries mistreated in any of these ways tend to have a very short working life indeed.

For a long time, sulphation was regarded as non-reversible and batteries that had lost too much capacity due to this effect were simply discarded. This was not only wasteful but was also an environmental problem, because both lead and sulphuric acid are highly toxic materials.

Around the middle of the last century, though, people in rural areas discovered that they could 'resuscitate' sulphated batteries by zapping them with high-voltage pulses from their electric fence controllers. They didn't exactly understand why this method worked but kept using it because it did.

Subsequently, in 1976, the US Patent Office granted a patent to William H. Clark of Salt Lake City, Utah, for a method of charging lead-acid batteries by means of narrow high-current pulses. This was claimed to more effectively dissolve the lead sulphate crystals and hence prolong battery life. Since then, a number of designs for pulse-type battery rejuvenators or

'zappers' have appeared in electronics magazines.

There is still a lot of argument about whether or not battery sulphation can be reversed and hence about the effectiveness of 'zapper' type pulse rejuvenators. Our prototype did initially seem to achieve a useful amount of rejuvenation on a badly sulphated battery (which later went short circuit) but we really cannot vouch for the overall effectiveness of this circuit. It simply hasn't been tested on a wide enough range of batteries.

However, it's cheap enough to build, so interested readers can put one together and try it out for themselves.

By the way, please note that there is evidence that only 'flooded' (liquid electrolyte) lead-acid batteries respond to this type of pulse desulphation. Sealed batteries with 'gel' electrolyte don't respond much at all, so we don't recommend using the zapper on this type of battery.

It's also worth noting that even on flooded lead-acid batteries, pulse desulphation is not quick. It can take tens or even hundreds of hours to achieve a significant amount of rejuvenation.

A problem with many of the published zapper designs, is that they use a P-channel power MOSFET. However, these are more expensive and harder to obtain than N-channel devices, so we've had quite a few requests for a design using one of the latter devices instead. And that's exactly what we've done, with the design described here using a low-cost IRF540N MOSFET.

How it works

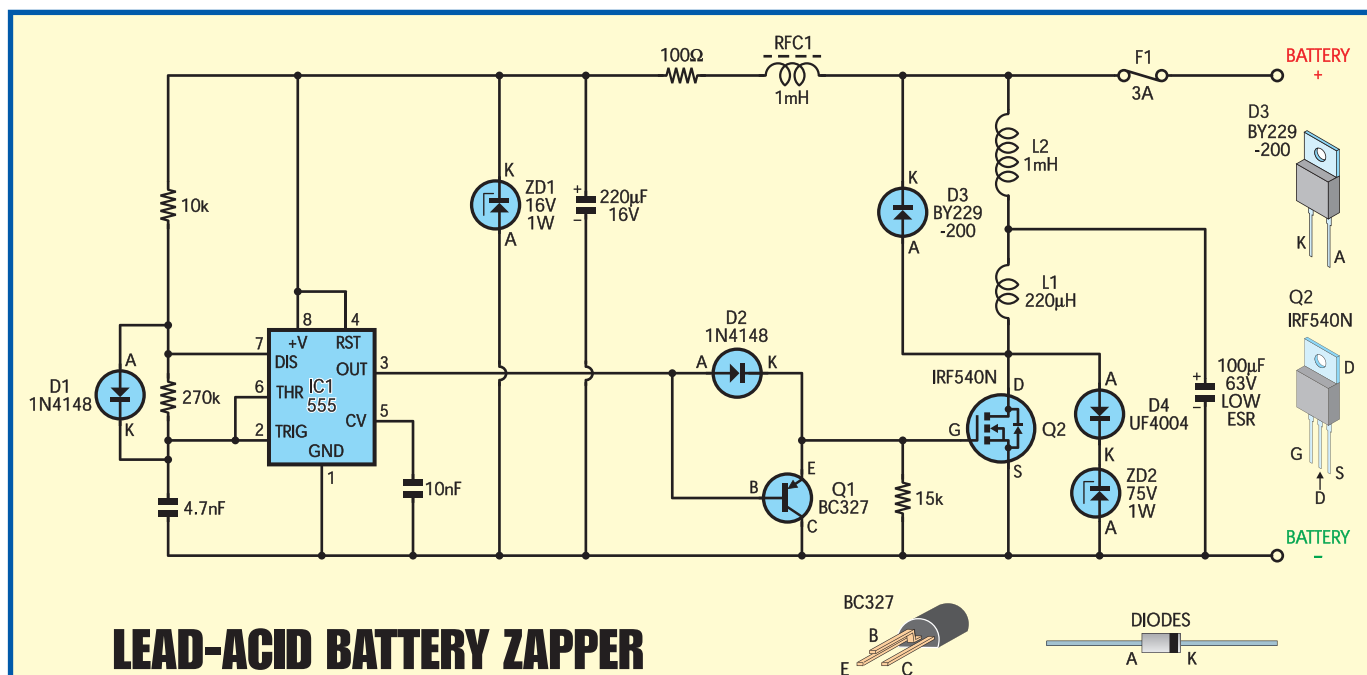
The basic principle used in desulphating zappers is quite simple: they draw a small amount of energy from either the battery itself or a charger connected to it, store this energy in a capacitor and then deliver it back to the battery as a narrow high-voltage pulse. In other words, a short pulse of current is forced through the battery in the 'charging' direction. It is these short current pulses which are claimed to dissolve the sulphate crystals (providing you're patient).

Disclaimer!

As stated in the article, our initial experiences with the Lead-Acid Battery Zapper indicated positive results. However, we must emphasise that our testing has been much too limited for us to make any claims or give any guarantees regarding the effectiveness of this unit.

In practice, you may find that the zapper successfully 'rejuvenates' some batteries, particularly if the battery has simply sulphated due to lack of use. However, it cannot possibly rejuvenate a battery that is worn out – ie, one in which the active material on the plates has been severely degraded.

Depending on the battery, it's also possible that any rejuvenation effects may be only temporary in nature.



LEAD-ACID BATTERY ZAPPER

Fig.2: the circuit for the battery zapper uses a 555 timer IC to turn MOSFET Q2 on for 50μs every 1ms (ie, at a 1kHz rate). Q1 shorts Q2's gate to ground each time IC1's pin 3 output switches low, to ensure a fast turn off.

Fig.1 shows the basic scheme. As shown, the circuit consists of two small inductors, a 100μF electrolytic capacitor, a fast-recovery diode (D3) and a high speed electronic switch. The switch is actually the N-channel power MOSFET (Q2) but it's shown in Fig.1 as a switch because that's how it's being used.

During the first phase of the circuit's operation (A), current flows from the battery (or charger) and charges the 100μF electrolytic capacitor via 1mH inductor L2. This charging phase lasts about 950μs, which is quite long compared with the next phase.

Next, during the second phase of operation (B), the switch (Q2) is closed. This connects 220μH inductor L1 to ground (battery negative), resulting in a sudden flow of current from the capacitor into L1. As a result, the energy stored in the capacitor is transferred to the inductor's magnetic field.

This phase only lasts for about 50μs – ie, just long enough for the energy transfer to take place.

At the end of the second phase, the switch is opened again (C). This sudden interruption of the inductor current causes an immediate reversal of the voltage across the inductor and


so a high-voltage pulse appears across the inductor, with the polarity shown. As a result, a discharge current pulse flows from the 100μF capacitor, down through L1, up through diode D3 and then out through the battery. This is the third phase of the circuit's operation.

This sequence of events is repeated indefinitely while the zapper is connected to a 12V battery (or battery and charger combination). That's because as soon as the discharge energy pulse from L1 has ended, the 100μF capacitor begins charging again via L2. So the remainder of the third phase becomes the first phase of the next charge-transfer-discharge cycle and that's how it keeps going.

Circuit details

Fig.2 shows the full circuit details of the Lead-Acid Battery Zapper. It incorporates all the circuitry shown in Fig.1, plus some extra parts to generate the short pulses to turn MOSFET Q2 on for 50μs every 1ms. In other words, Q2's gate is driven with 50μs-wide positive pulses at a rate of 1kHz, which means that the pulses are spaced 950μs apart.

This train of narrow pulses is generated by 555 timer IC1, which is connected as an astable oscillator. Diode D1, the 10kΩ and 270kΩ resistors, and the 4.7nF timing




capacitor ensure a very high mark-space ratio at the pin 3 output. In operation, D1 ensures that the 4.7nF capacitor charges up very quickly via the 10k Ω resistor but can only discharge relatively slowly via the 270k Ω resistor (ie, when the internal discharge transistor, on pin 7 turns on). As a result, IC1's pin 3 output goes high for 50 μ s, then low for 950 μ s and so on.

Transistor Q1 and diode D2 are used to ensure that the pulse stream from pin 3 of IC1 turns switch Q2 on and (especially) off very rapidly. In effect, they compensate for the charge stored in Q2's gate-channel capacitance when the MOSFET is turned on.

They do this very simply: when IC1's output goes high, D2 conducts and the pulse is applied directly to Q2's gate to turn it on. When IC1's output subsequently drops low again, this suddenly turns on transistor Q1 and effectively connects a short-circuit between Q2's gate and ground. As a result, the gate charge in Q2 is discharged very rapidly, making Q2 turn off again in very short order.

There's very little else left to explain. Inductor RFC1, the 100 Ω series resistor and Zener diode ZD1 allow the +12V DC rail to be applied to IC1 but block the high-voltage pulses generated in the output stage from reaching the IC. Fuse F1 is there to protect the circuit



The PC board fits neatly inside a standard utility box and is secured using 6mm spacers and machine screws and nuts.

from damage if the supply leads to the battery (or charger) are connected with reverse polarity.

Finally, D4 and Zener ZD2 form a clamp circuit to protect MOSFET Q2 from voltage spikes.

Construction

Construction of the Lead-Acid Battery Zapper is straightforward, with all parts (except for the fuse) mounted on a PC board, coded 623 and measuring 122 x 57mm. This board has cutouts in each corner so that it fits snugly

inside a standard utility box (130 x 67 x 44mm).

Fig.3 shows the assembly details. As usual, it's easiest to fit the low profile resistors and inductor RFC1 first, followed by the smaller capacitors and then the electrolytics. Note that the electrolytics are polarised, so make sure they go in the right way round.

Next, fit diodes D1 and D2, again taking care to ensure correct polarity. The same applies to Zener diode ZD1, which can also now go in.




Fig.3: follow this parts layout diagram to build the PC board.

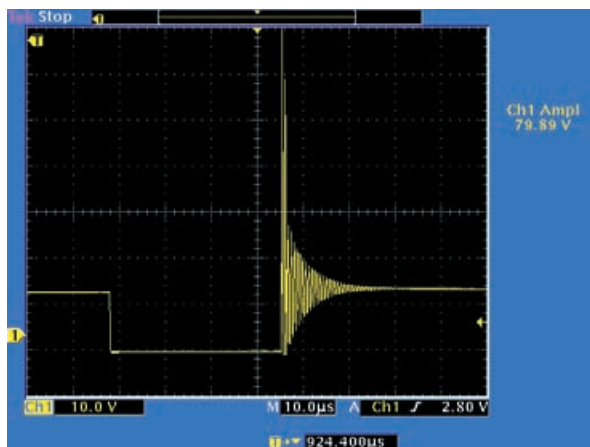


Fig.4: this scope shot shows the pulse waveform at the drain of MOSFET Q2. Note the ringing in the pulse waveform following the main voltage spike.

Reproduced by arrangement with
SILICON CHIP
magazine 2007.
www.siliconchip.com.au

That done, fit transistor Q1, MOSFET Q2 and diode D3, which is in a 2-pin TO220-style package similar to the package for Q2. These devices are all polarity sensitive, so again follow Fig.3 carefully to ensure correct orientation. Follow these parts with IC1, which should be fitted with its notched end towards the 270kΩ resistor.

The last components to fit are the two large air-cored inductors (L1 & L2). These are wound on plastic bobbins, with their wire ends emerging from holes or slots in the lower cheek.

Securing the inductors

Both inductors on the prototype were secured to the board using Nylon spacers inside their centre void, with a screw at each end, along with an M3 flat washer and 16mm grommet at the top of L1. This is the method shown in the photos and on the wiring diagram (Fig.3).

Note that, in each case, the inductor's leads must be passed through their matching holes in the PC board before they are secured in position. Once they're in position, the assembly is turned over and their leads soldered to their board pads.

The PC board assembly is now complete. However, before fitting it into the box, it's a good idea to solder the two supply leads to their pads at the righthand end of the board. Just

strip 4mm of insulation from the end of each length of cable, pass these down through their respective holes in the PC board (red to positive, black to negative) and solder them to the PC pads underneath.

Final assembly

The PC board assembly is supported inside the case on four M3 x 6mm untapped spacers and secured using M3 x 12mm countersink head screws, lockwashers and nuts.

The first step is to use the board itself as a template to mark out the

mounting holes. That done, remove the board, drill the holes to 3mm, and use an oversize drill-bit to countersink the holes from the back of the case.

A further two holes are required at one end of the case to pass the battery leads and these can be drilled to 4mm about 10mm down from the top. The panel-mount fuseholder is mounted at the other end of the case and requires a shaped hole to suit the threaded body. This hole can initially be drilled to 4mm, then carefully enlarged using a tapered reamer and shaped using a small flat file.

That done, the board assembly can be fitted to the case. This is done by first installing the four screws and fitting the 6mm-long spacers, after which the board assembly can be lowered into position while feeding its negative (black) power lead out through its matching hole at one end. It's then simply a matter of fitting the lockwashers and nuts and tightening up the screws, to secure the assembly in place.

The next step is to cut the positive (red) input/output lead about 120mm from the end of the board and remove about 5mm of insulation from the free end. That done, fit the fuseholder to the lefthand end of the case, with its side solder lug uppermost for access, and solder the positive lead from the PC board to it.

The remaining red lead can then be passed through its hole in the case and soldered to the fuseholder's other lug. Note that you will have to dress this lead carefully around L2 and the upper tabs of D3 and Q2, so that it reaches the fuseholder without strain.

Finally, complete the construction by fitting the lid to the case and attaching the two 32mm alligator clips to the far ends of the two input/output leads. Be sure to fit the red clip to the positive lead and the black clip to the negative lead. Your battery zapper is now complete and ready to use.

WARNING!

Hydrogen gas (which is explosive) is generated by lead-acid batteries during charging. For this reason, be sure to always charge batteries in a well-ventilated area.

Never connect high-current loads directly to a battery's terminals. Similarly, when using a battery charger, always connect its output leads to the battery before switching on mains power. Failure to observe these simple precautions can lead to arcing at the battery terminals and could even cause the battery to explode!

Note too, that the electrolyte inside lead-acid batteries is corrosive, so wearing safety glasses is always a good idea.

Table 1: Resistor Colour Codes

	No.	Value	4-Band Code (1%)	5-Band Code (1%)
□	1	270kΩ	red violet yellow brown	red violet black orange brown
□	1	15kΩ	brown green orange brown	brown green black red brown
□	1	10kΩ	brown black orange brown	brown black black red brown
□	1	100Ω	brown black brown brown	brown black black black brown

Fitting An On/Off Switch

Although not fitted to the prototype, we strongly recommend that a switch be installed in series with the positive battery lead to allow the unit to be isolated during connection and disconnection. This eliminates the possibility of arcing at the battery terminals.

Any miniature mains-rated switch would be suitable. It can be mounted on one end of the case, next to the fuse.




Fig.5: how to install the on/off switch. The 10nF capacitor across the switch reduces contact arcing.

A 10nF 100V polyester capacitor must be fitted directly across the switch terminals, as shown in Fig.5.

Parts List – Lead-Acid Battery Zapper

- 1 PC board, code 623 available from the *EPE PCB Service*, size 122 x 57mm
- 1 utility box (130 x 67 x 44mm)
- 4 6mm-long untapped metal spacers
- 4 M3 x 12mm machine screws, countersink head
- 4 M3 nuts and star lockwashers
- 1 220µH air-cored crossover inductor (L1)
- 1 1mH air-cored crossover inductor (L2)
- 1 1mH RF choke (RFC1)
- 4 plastic cable ties (to secure inductors L1 & L2)
- 1 M205 panel-mount fuseholder
- 1 3A slow-blow M205 fuse
- 1 1.5-metre length of heavy-duty cable, red insulation
- 1 1-metre length of heavy-duty cable, black insulation
- 1 pair of 32mm alligator clips (red and black)

Semiconductors

- 1 555 timer (IC1)
- 1 BC327 PNP transistor (Q1)
- 1 IRF540N N-channel 100V/12A MOSFET (Q2)
- 1 16V 1W Zener diode (ZD1)
- 1 75V 1W Zener diode (ZD2)
- 2 1N4148 diodes (D1,D2)
- 1 BY229-200 fast-recovery diode (D3)
- 1 UF4004 ultra-fast diode (D4)

Capacitors

- 1 220µF 16V radial elect.
- 1 100µF 63V low-ESR radial electrolytic
- 1 10nF polyester
- 1 4.7nF polyester

Resistors (0.25W 1%)

- 1 270kΩ 1 10kΩ
- 1 15kΩ 1 100Ω

remove the power and write that battery off as one that cannot be saved. In other words, there are no guarantees that the zapper can resurrect *all* badly sulphated batteries – it can't. **EPE**




Fig.6: here's how to use a charger with the Battery Zapper. Note the 1mH inductor in series with the charger.

Putting it to use

Using the zapper is easy – just connect its leads to the terminals of the battery you want to rejuvenate (red to positive, black to negative).

There's only one qualification: if the battery is already so discharged that it can't supply the 50mA or so needed to operate the zapper, you'll need to connect a conventional trickle (or low-current) charger to the battery as well – at least to get the rejuvenation process started (see Fig.6). And if the battery is very badly sulphated as well, you'll have to keep the charger connected for quite a while.

After that, it's simply a matter of leaving it to pulse away until the sulphate crystals inside the battery have dissolved. This can take quite some time – from a few days to a few weeks – so you need to be patient.

If your charger doesn't have an in-built current meter, you can connect an ammeter in series with one of its leads so that you can monitor the charging rate. This should increase slowly as the sulphate crystals dissolve.

By the way, if you do have to connect a charger to the battery to power the zapper, you *must* use a

WARNING!

This circuit generates high-voltage pulses which could easily damage the electronics in a vehicle. *DO NOT* connect it to a car battery installed in a vehicle.

1mH air-cored inductor (the same as L2) in series with one of the charger's leads (see Fig.6). There are two reasons for this: (1) to protect the output circuitry of the charger from possible damage; and (2) to prevent the charger's relatively low output impedance from shunting the pulses, thereby reducing their effectiveness.

It doesn't always work

A final warning: not all lead-acid batteries are capable of being desulphated by this zapper. In some batteries, the lead-sulphate crystals stubbornly resist the pulsing effect and the battery can sometimes even develop a short-circuit between the plates.

So, if the battery charger current suddenly increases to a very high level,

All about PIC Special Features

It's back to datasheets this month, this time looking at the 'loose ends' that Microchip wrap up in a chapter called 'Special Features of the CPU'. Once again, the datasheets for all the 8-bit series of PIC processors follow a common format and they all contain this chapter, with more or less of the features depending on the processor variant.

Issues covered in 'Special features' do not easily fit into a normal peripheral feature chapter, but they are vital features that can have a significant effect on the operation of the processor. These features are configured by programming 'configuration bits', and the default settings of many of them will cause the processor to not execute your program. So there are two aspects to getting your program to work: writing the software, and choosing the appropriate settings for the configuration bits.

Configuration bits

The configuration bits are the most important issue covered in this section of the datasheet. These are non-volatile 8-bit registers inside the processor that configure various features of the device. Being non-volatile, they maintain their values when power is removed, just like code memory (these registers are in fact stored in flash memory, but in an area outside of the normal program memory address range.) In most cases you would never access the configuration bits in your application, but instead set them when the application software is being programmed into the device.

There are essentially two areas of configuration memory: Config bits and Device ID bits. To someone new to the PIC architecture the storage of these bits can be somewhat confusing. Fig.1 details the memory layout.

The PIC is a Harvard Architecture processor, which means the flash and RAM memory areas have their own independent address and data busses. This explains why some of the lower address ranges are duplicated. The processor understands from the context of each instruction word being executed which type of memory to access. Some access data memory, some program memory.

The Config bits reside in an area of non-volatile memory starting at address 300000h. Note that this does not mean you have over 3MB of memory! The top four bits of the address, '3', tells the processor to switch between code memory and configuration bit memory, and then the lower bits of the address are used as an index into

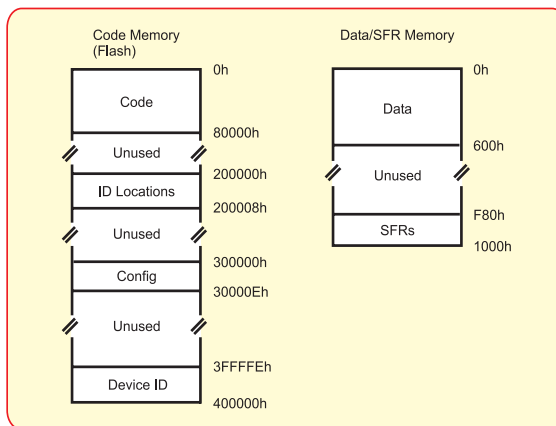


Fig.1. PIC memory layout

the selected memory area. The address of 300000h was probably chosen by Microchip to allow space for the on-chip flash memory to be expanded in newer versions of processors, without having to change the addresses of the configuration registers. 3MB should be plenty of space!

If you study the tables within the datasheet that define the locations of the Config bits you may notice that some of the registers are not in contiguous locations. This is because Microchip often allocates specific config register addresses to specific peripheral hardware features. If a certain peripheral feature is not present on your particular device then the config registers associated with it will be absent and the relevant addresses unused. This is just another example of Microchip's smart approach to developing processors: re-using parts of designs throughout the processor range.

As we mentioned earlier, the Config bits are split into two distinct areas. The Config bit registers provide the options for all the different peripheral features of the CPU. The Device ID registers are read-only bits and serve to identify the type of device the part is. It's the Device ID bits that enable the programmer software to automatically detect what chip is inserted in the programmer. The bits are also accessible by your own application software – perhaps your program might check that it is running on the correct type of PIC!

The Config bits are writable, and it is up to you to specify the values they should be set to, typically in one of your source files. It's advisable to set all the Config bits, even the ones whose default values are acceptable, since the default setting may change in the future. You must study the definition of

these Config bits, understand them and write in the correct values. The default settings of the Config bits are almost certainly not going to work with your software; you will have to change at least some of them, so take the time to understand what they all do.

ID locations

Hidden in the detail of the Config bit memory locations are a number of locations reserved for 'ID locations'. These are a small number of non-volatile bytes that you can use to store information

about your application, such as a version number and identifier.

These locations can be read even when the code memory itself has been read-protected. This can be very useful to identify what is inside a read-protected device at a later date. So, if you intend to read-protect your code we recommend you use the ID locations.

Config registers

There are a number of ways in which you can specify the contents of the config registers, and the exact mechanism will be dependent on which compiler, assembler or programmer software you use. The normal way is to specify them in your source code files. When the source file is assembled or compiled, the settings get copied into your programming file and transferred over by your programming software.

You don't have to do this, however; most programming software allows you to view, edit and change the contents of the configuration bits by hand. But you don't really want to do this – it makes much more sense to place them in your source files so they will not get lost in the future.

In MPLAB assembler you can use the CONFIG directive to set individual bits, like this:

CONFIG WDT=ON, WDTPS=128

There is also another, older directive that can be used, `_config`, which allows you to specify the address of the register and its content directly:

_config 0x300002, 0x18

This method is deprecated now and CONFIG is the preferred method since it doesn't tie features to specific address locations.

Setting config registers in the Microchip C compiler is just as simple, using the `#pragma` directive:

```
#pragma config WDT = ON, WDTPS  
= 128
```

You only need to specify the config settings once in your source files, typically at the beginning of the main file.

Other assemblers and compilers will have slightly different syntax but should be similar.

The Config bits specify the operation of the various 'special features' of the CPU, and as they are all located within the same area of memory, Microchip have chosen to describe them all within the same chapter, with some additional information being found elsewhere in the datasheet without a cross reference or an obvious entry in the index. This makes for a very concise and sometimes confusing description, which is a pity as the correct setting of these bits is vital to the proper operation of the device. (It's a small criticism; in the author's view Microchip produce the most comprehensible processor datasheets on the planet!)

To help explain these features we will go through the Config bits of an example processor, the PIC18F2520, a device with a wide range of features shared by many of the PIC range. We will walk through each config register in turn, starting with CONFIG1H.

Clock source

The four least significant bits, FOSCO – FOSC3, are probably the most important ones as they define the source of the main processor clock, the signal that drives the execution of the program, and pretty much everything else. Fortunately, the options are clearly defined in section 2 of the datasheet. If you want to run from a standard crystal at a frequency of 4MHz or higher (which seems to cover most projects), you need to select the HS option, which corresponds to FOSC bit settings of 0010. You can specify this in your main source code file using the directive:

```
CONFIG OSC=HS
```

These special CONFIG setting values (OSC in the example above) can be found in a Microchip document called 'PIC18 Configuration Settings Addendum DS51537F' which can be downloaded from the Microchip website (www.microchip.com).

Back to the CONFIG1H, the IESO bit enables the oscillator switch over feature. Crystal oscillators can take several milliseconds to reach full oscillation following power-up, which can, in rare cases, be an inconvenience. The IESO feature enables the processor to start running immediately off the internal RC oscillator and then automatically switch over to the crystal oscillator once it has stabilised. It's not an essential feature and best left disabled. For the curious, it is described in section 2 of the datasheet.

The FCMEN bit controls the Fail safe Clock Monitor feature. This enables the

CPU to monitor the main external oscillator and switch over to the internal RC oscillator if it should stop for any reason. Without this your software will simply hang in the unlikely event that the oscillator should stop. A very useful feature should you require ultimate reliability and safe operation, but not really necessary for hobby applications. The feature is described in more detail later in the Special Features chapter.

Brown-out Reset and Power-up Timer

The register CONFIG1L is not implemented in this processor, so the next register is CONFIG2L. This provides control of two independent features: Brown-out Reset control and the Power-up Timer. The Power-up timer is a simple counter register, clocked from the internal RC oscillator, that will delay the start of the CPU by approximately 65ms following the release of the reset signal. We strongly advise users to enable this feature as it allows the external crystal oscillator time to stabilise before the CPU starts using it. External oscillators can have unpredictable behavior while powering up, giving out clock pulses that exceed the specification of the PIC – which can result in the CPU getting into an odd state that it cannot recover from.

The Brown-out Reset control is also an important feature that one should consider enabling. 'Brown-out' refers to the condition where the supply voltage 'dips' briefly below a safe operating voltage, but is high enough (and short enough in duration) to not cause a full device reset. These brown-outs can cause erratic software operation, and are not unusual – especially during power-up and power-down. Software that accesses the Data EEPROM during these times is particularly sensitive to erratic supply rail behavior, and it is not uncommon to see EEPROMs corrupted under these conditions.

The Brown-out reset feature enables the close monitoring of the supply rails against a user-selectable level, and forces the processor into the reset state when the voltage dips below it. The BORV0 – BORV1 bits enable you to select one of four levels, and it is a case of experimenting with them to see which provides the most reliable operation. If your power supply is going to be a battery, or is not going to provide a fast switching stable voltage, then you should enable this feature, using the BOREN0 – BOREN1 bits. Brown-out operation is detailed further in section 4 of the datasheet.

Watchdog Timer

CONFIG2H controls the Watchdog Timer function. This is a simple timer that is clocked by an on-chip oscillator, and so is independent of the main CPU oscillator. The idea behind the Watchdog is that you should periodically issue a CLRWDT instruction to set the timer back to zero. If your software fails to issue the instruction within a certain time (the timeout time) then the processor will reset, starting the

application from the beginning. Watchdogs are a simple and effective means of catching unusual 'lock-up' problems, although they should not be used as a substitute for good design!

The Watchdog timeout can be set using bits WDTPS3 – WDTPS0, giving a range of timeout times from 4ms to over 128s. It's a very useful feature but you have to think carefully about where you place your CLRWDT instructions. Accidentally enabling this feature before you have written the code to reset the timer is a common cause of those unusual bugs where the code looks perfect but simply doesn't work, or only runs for a few seconds. Keep the feature disabled until your software is working, and then think about adding it in afterwards for extra protection.

Misc Control Flags

CONFIG3H provides an assortment of single bit control flags. MCLRE allows for the main reset input signal to be disconnected from its pin, freeing up the pin to become an additional I/O called RE3.

Bringing the reset signal inside the chip removes the flexibility of being able to set your own reset time (using an RC circuit connected to the pin). Unless you are desperate for the additional I/O pin, keep the reset signal connected to the reset pin by setting the MCLRE config bit to 1.

Timer1

Timer1 has the ability to drive another external crystal oscillator, typically a low power 'watch crystal' for providing a real time clock source. The Config bit LPT1OSC enables the user to set the power level with which this oscillator is driven. It should normally be left set to 0, but by setting it to 1 you can reduce the device current consumption. Low power operation can be susceptible to noise in the circuit, so it's best to leave this bit set to 0 unless you are confident in the design of your circuit and its operating environment.

PortB I/O Pins

Bit PBADEN defines to which peripheral the PORTB0-PORTB1 I/O pins are initially connected to following a reset. When this bit is set to 0, the pins are digital I/O. When set to 1, the pins are connected to the Analogue-to-Digital converter. The SFR ADCON1 can be used to change the assignments afterwards; this bit is provided to enable the correct setting to be applied immediately following reset and before any code is run.

Bit CCP2MX allows you to specify to which I/O pin the CCP2 I/O signal is routed; When set to 0 the signal is routed to RB3, when set to 1, RC1. This can be very useful for example, if you need to implement an 8-bit port on PORTB, but still want to use the CCP2 signal. This type of option is often found on the larger microcontrollers that provide more I/O signals than they have pins. Although it gives you more to think about, it provides for greater flexibility in determining how you connect to the PIC.

DEBUG and XINST

CONFIG4L also contains an assortment of control bits, including a very important one. The DEBUG bit is used by debugger hardware like the PICKit2, and you should always leave this programmed to 0 in your application. The XINST bit is used to switch between normal and extended instruction set. The extended instruction set is provided to make programs written in 'C' more efficient; compilers like Microchip's MCC18 have support for operating in either normal or extended mode. The operation of the processor in extended mode is somewhat more complicated and so if you are writing software in assembly language it is best to leave this bit set to 0 (normal or legacy mode) unless you like a challenge!

Stack Overflow

The STVREN bit enables a new feature: Reset the device on stack overflow. The stack is only 31 levels deep, which means you can only have up to 31 levels of nesting within your software. For example, if your main routine calls a sub-routine A, then the code in sub-routine A is running at a nesting level of 1. If sub-routine A itself calls another routine B, then code in that routine is running at a nesting level of 2, until it performs a return to sub-routine A, when the nesting level returns to 1.

This nesting can quickly mount up if you organise your code in sub-routines (which is a good idea). Under normal conditions, should the stack overflow then you will find your code operating erratically, returning to the wrong routines. Enabling the stack overflow reset feature will cause the device to do a full reset should this overflow occur, which will hopefully be very easy to detect when you are testing your code. It's a useful feature and worth enabling.

Low voltage programming

The last bit in CONFIG4L is the LVP or low voltage programming bit. This feature is present in many PIC variants and often causes problems due to the way in which it works. To understand why, a bit of background. The normal way to put a PIC into programming mode is to raise the MCLR pin to +12V. By bringing the pin to a high, non-standard level, it signals to the CPU that entry into programming mode is

required. It's not always convenient to provide +12V, so Microchip have provided another way to bring the CPU into programming mode.

By setting the LVP bit (using a programmer equipped with the standard, 12V programming signal), the microcontroller can then subsequently be reprogrammed at 5V using the PORTB.5 pin. Doing so means that PORTB.5 can no longer be used as an I/O pin, and more importantly, must be tied to V_{SS} during normal operation – failure to do so will mean your application will fail to start.

Other Config Registers

The remaining Config registers are all related to protecting access to various parts of the memory. Keeping these bits set to 1 marks all areas of memory as unprotected. In this state anyone using a simple programmer unit can read the contents of your device – which would be an issue if you want to release products without others being able to copy the contents. By clearing these Config bits you mark areas of memory as un-readable, protecting your software from prying eyes. The Config bits give you a very fine granularity of control over which sections of memory can be accessed.

Areas of memory that are read-protected have to be erased to regain access to them. As mentioned earlier, the ID locations within the Configuration Memory space are always readable, which enables you to program some kind of software identifier into them – so you can work out what is inside a read-protected chip at a later date, should you forget!

In summary

Some of these special features can be difficult to get to grips with. The best approach to overcoming this is to work out the minimum essential features (such as the oscillator mode) and leave the other functions disabled until you have code running on your board.

Once you are happy that the basic design is working, start introducing additional special features one by one, debugging and understanding each one before moving on to another. After a few projects you will find them easy to use, and as the features are often common across different processors you will find learning another processor architecture straightforward. Just take it a step at a time!

Reset values

To finish of this month we will look at another interesting feature of the processor, the default values that Data Memory and Special Function Registers get set to following power-up or on reset. Many of the SFRs and all of the Data Memory is left in an undefined state, which means sometimes the bits will be zero, and sometimes they will be ones. If the processor goes through a reset without power being removed from the device then data memory locations will maintain the values they had before the reset – a feature that can be usefully exploited in some circumstances.

Some SFRs will be set to specific values after a reset, since without doing so the operation of the processor would be unpredictable. You can see what the reset values of an SFR will be by looking at the section in the datasheet which defines it. Each bit in an SFR is shown with its operation mode and reset value, for example:

R/W-0

means the bit can be read, written to and has a value after reset of 0.

The datasheet will contain a more comprehensive list of default settings for SFRs, called 'Reset State of Registers'. This shows the state of each SFR under different reset conditions – yes, the different causes of reset can leave your SFRs in different states.

As Data Memory can have unpredictable values in it following a power-on reset it is always advisable to clear all your data memory to a known value (FF or 00) on power up. It is not uncommon to find a program that works perfectly until you leave the power disconnected for an extended period, placing a random value into a critical (but uninitialised!) variable. These kind of problems can be very hard to solve.

Clearing memory removes this problem and is fast and easy: the following code segment clears all bytes in bank 1:

```
lfsr   FSR0, 100h
mc001:
clrf   POSTINC0
btfss  FSR0H, 1
bra    mc001
```

Discovering and then tracking down bugs caused by uninitialised variables can be very time consuming so unless your code is very simple, the technique above is well worth adding.

EPE EVERYDAY PRACTICAL ELECTRONICS

NEWSAGENTS ORDER FORM

Please reserve/deliver a copy of *Everyday Practical Electronics* for me each month

Name and Address

Postcode Tel

Everyday Practical Electronics is published on the second Thursday of each month and distributed S.O.R. by SEYMOUR
Make sure of your copy each month – cut out or photocopy this form, fill it in and hand it to your newsagent

By JIM ROWE

The on-screen video looks considerably better than this photo indicates. The contrast is better and the Moire patterning, a result of the interaction between the on-screen display and our digital camera, is absent.

Video Reading Aid

...for vision impaired people

Do you have a family member with vision problems – like cataracts, or age-related macular degeneration? Here's a low-cost video reading aid that will make it much easier for them to read a book or newspaper. It combines a small CMOS TV camera with a video processor which boosts the contrast and allows them to select either a positive or negative enlarged image for viewing on a TV set or video monitor.

EYE PROBLEMS like cataracts and age-related macular degeneration are all too common, especially among those of 'mature age'. In fact, it was recently estimated that one in every four people over 75 has symptoms of this kind of visual impairment, while one in every 10 lose their central vision.

Understandably, those unlucky enough to suffer from these problems can find it very difficult to read a book, magazine or newspaper. This lowers their quality of life dramatically and deprives them of important

sources of news, entertainment and information.

In many cases, however, reading printed material can be made a lot easier by using improved lighting to increase the contrast, plus a magnification system to enlarge the type. Optical magnifiers with built-in lighting are available for use as reading aids but they're fairly pricey. You can also get similar devices using video magnification but these are even more expensive. As a result, such devices are often out of the reach of the people who could benefit from them.

Recently, we decided to have a go at a video magnifier ourselves and this project is the result. It combines one of the very small low-cost black and white CMOS cameras currently available from various suppliers with a very compact video processing circuit, and has a switch so you can select one of three image options: high contrast greyscale positive, hard limited or 'digital' black and white positive, or digital negative. And the output is standard video so it's compatible with any normal PAL TV receiver.

The camera and video processor are both fitted inside a standard project box. Because a person with impaired vision doesn't want to be fiddling with camera focusing, we've mounted it on a plastic food container to give it a fixed focal length. In use, this plastic skirt sits directly on the printed page and slides easily over the page, without marking.

Basically, it behaves a bit like a giant mouse – you just slide it so that the lens is over the text you want to read.

Illumination is provided via four high-output white LEDs, which mount on the underside of the box adjacent to the lens. In practice, the LEDs have to be 'doctored' to ensure that their light output is reasonably diffused over the camera's viewing area but this is easy to do, as described later in the article.

The end result is an easy-to-build video magnifier which you can feed into almost any old colour or B&W TV set.

The design uses one of the low-cost B&W cameras with a CMOS sensor that are currently available from various electronics retailers. We've tried it out using two of these: the Swann unit and the Samsung unit. These both give good results, although the Swann unit requires a minor modification to disable its inbuilt IR LEDs, so that it runs cooler (more on this later).

Of course, other mini CMOS cameras should also be suitable.

How it works

Refer now to Fig.1 for the circuit details. The output of the CMOS camera is fed through a video processing circuit that's rather similar to some video enhancers but modified to enhance the contrast. The circuit can also generate a negative version of the image, without degrading the signal's sync pulses.

As shown, the video output from the camera is terminated in a 100Ω load, to provide matching. It then passes through a $1\mu\text{F}$ coupling capacitor, after which it splits in three directions: across to CMOS analogue switch IC2a, down to the pin 2 input of sync separator chip IC4 (via a 100Ω resistor and a 100nF capacitor) and further down to the non-inverting (pin 3) input of video amplifier stage IC5a.

IC4 (the sync separator) is used to extract the sync and 'back porch'



The Video Reading Aid skates over the printed page on a plastic skirt (actually an upside down food container). This keeps the lens at the correct focal distance and makes the unit easy to operate.

gating pulses from the video signal. These are then used to provide control signals for video switches IC2a and IC2b.

In greater detail, both the back porch and composite signals are combined in gate IC3c (used here as a negative-input OR gate) and used to turn on switch IC2a, to allow the sync and blanking information to pass straight through. At the same time, IC3a inverts this signal to control switch IC2b. This latter switch allows the processed video through to the output buffer (IC5b) during the 'active' part of each video line.

In effect, IC2a and IC2b operate in complementary fashion. When IC2a is on (closed), IC2b is off (open) and vice versa. This means that when IC2a is closed, the sync and blanking pulses are fed through to IC5b while the active video is blocked. Conversely, when IC2b is closed, the active video is fed through and the sync signal is blocked.

The 'back porch' (or burst gating) pulses from pin 5 of IC4 are also inverted by IC3b and used to control switch IC2c. This forms an active clamp to fix the blanking level of the incoming video to ground potential.

The part of the circuit we've just described is basically the control section, which ensures that only the active video is subjected to processing.



The Video Reading Aid is based on a miniature black and white CMOS camera, such as this Swann unit.

Now let's look at the actual processing circuitry itself, which involves IC5a, IC6, transistor Q1 and IC2d.

IC5a is simply a video amplifier and operates with a fixed gain of two, as set by the two 510Ω resistors in its feedback divider. Its output at pin 1 becomes the 'high contrast analogue positive' video signal and is fed to the first position of selector switch S1.

This same output signal is also fed to the non-inverting input (pin 2) of IC6, an LM311 high-speed comparator. This compares it with a reference DC voltage level on pin 6, as set by trimpot VR1, to generate a 'hard limited' or rectangular digital equivalent of the boosted video signal.

IC6 has positive feedback applied via the $4.7\text{k}\Omega$, 100Ω and $33\text{k}\Omega$ resistors, to

Parts List

- 1 PC board, code 624, available from the *EPE PCB Service*, size 122.5 x 57.5mm
- 1 utility box, 130 x 67 x 44mm
- 1 mini CMOS B&W TV camera (see text)
- 2 L-brackets, 15 x 15 x 10mm – see text
- 1 47 μ H RF choke (RFC1)
- 1 3-pole 4-way rotary switch (S1)
- 1 2.5mm PC board mounting DC connector (CON1)
- 1 RCA phono connector, PC board mounting (CON2)
- 1 4-pin SIL header strip
- 4 M3 x 25mm tapped metal spacers
- 8 M3 x 6mm machine screws
- 2 M3 x 10mm machine screws
- 1 3m length of light figure-8 twin shielded cable
- 2 RCA phono plugs, yellow
- 1 12V/200mA regulated plug-pack supply, with 2.1mm plug
- 1 2.1mm concentric DC line socket (to match plugpack)
- 1 plastic food container, 130 x 105 x 60mm
- 1 1k Ω mini horizontal trimpot (VR1)
- 1 5k Ω mini horizontal trimpot (VR2)

Semiconductors

- 1 741 op amp (IC1)
- 1 4066B quad bilateral switch (IC2)

- 1 4093B quad CMOS Schmitt NAND gate (IC3)
- 1 LM1881 video sync separator (IC4)
- 1 MAX4451ESA dual video amp (IC5)
- 1 LM311 comparator (IC6)
- 1 PN100 NPN transistor or BC548 (Q1)
- 4 5mm high-brightness white LEDs (LED1-LED4)
- 3 1N4148 signal diodes (D1,D2,D5)
- 2 1N4004 power diodes (D3,D4)
- 1 1N752 5.6V/400mW Zener diode (ZD1)

Capacitors

- 1 220 μ F 16V radial elect.
- 1 10 μ F 10V radial elect.
- 3 4.7 μ F 16V tantalum
- 1 1.0 μ F MKT metallised polyester
- 2 100nF MKT metallised polyester
- 6 100nF multilayer monolithic
- 1 2.2nF 50V disc ceramic
- 1 220pF 50V disc ceramic
- 1 22pF 50V disc ceramic

Resistors (0.25W, 1%)

- | | |
|-----------------|-----------------|
| 1 680k Ω | 1 2.2k Ω |
| 1 33k Ω | 2 1k Ω |
| 1 4.7k Ω | 4 510 Ω |
| 1 3.9k Ω | 4 270 Ω |
| 1 3.3k Ω | 4 100 Ω |
| 1 2.7k Ω | 1 75 Ω |

through video buffer IC5b, which operates with a fixed gain of two, to compensate for the losses in the 75 Ω ‘back termination’ resistor in series with the output. This is the standard video buffer configuration and is used to allow the output signal to be fed along relatively long video cables with minimal degradation.

Power supply

Both the mini video camera and the video processing circuitry are powered from an external 12V DC source – either a 12V battery or a regulated plugpack supply, delivering 12V at up to about 150mA. The four white LEDs (LED1-LED4), used to provide illumination, are powered from the same source.

Series diodes D4 and D5 provide reverse polarity protection and also reduce the overall supply voltage to 10.8V, which is necessary to protect both IC5 and the CMOS camera from over-voltage damage. Because IC5 needs a balanced bipolar supply, IC1 and Zener diode ZD1 are used to give the 10.8V supply an active ‘centre tap’, which is connected to the circuit’s earth. The two main supply rails thus become +5.4V and -5.4V nominal with respect to ground.

The CMOS camera and all of the remaining ICs are connected directly between the +5.4V and -5.4V rails, as are the illumination LEDs. The latter are connected in two series strings, with a 270 Ω resistor in each string to limit the current to around 17mA. Provided high-brightness white LEDs are used, this modest current provides plenty of illumination.

Construction

All of the video processing circuitry fits on a PC board measuring 122.5 x 57.5mm and coded 624. This board has a rounded cutout in each corner, so that it slips neatly inside a standard utility box. The video selector switch is located near the centre of the board, while the DC input and video output connectors are mounted at one end – see Fig.2.

The CMOS camera module is mounted centrally inside the box (Fig.6). It sits under the PC board with its lens protruding through a 16mm hole in the box base and is supported by two small aluminium angle brackets. The adjacent illumination LEDs are mounted on the copper side of the PC board at

give it a small amount of hysteresis and ensure clean switching. Trimpot VR2 also allows fine adjustment of this feedback. The output from pin 7 is then fed to transistor Q1, which is connected as an emitter follower to provide buffering.

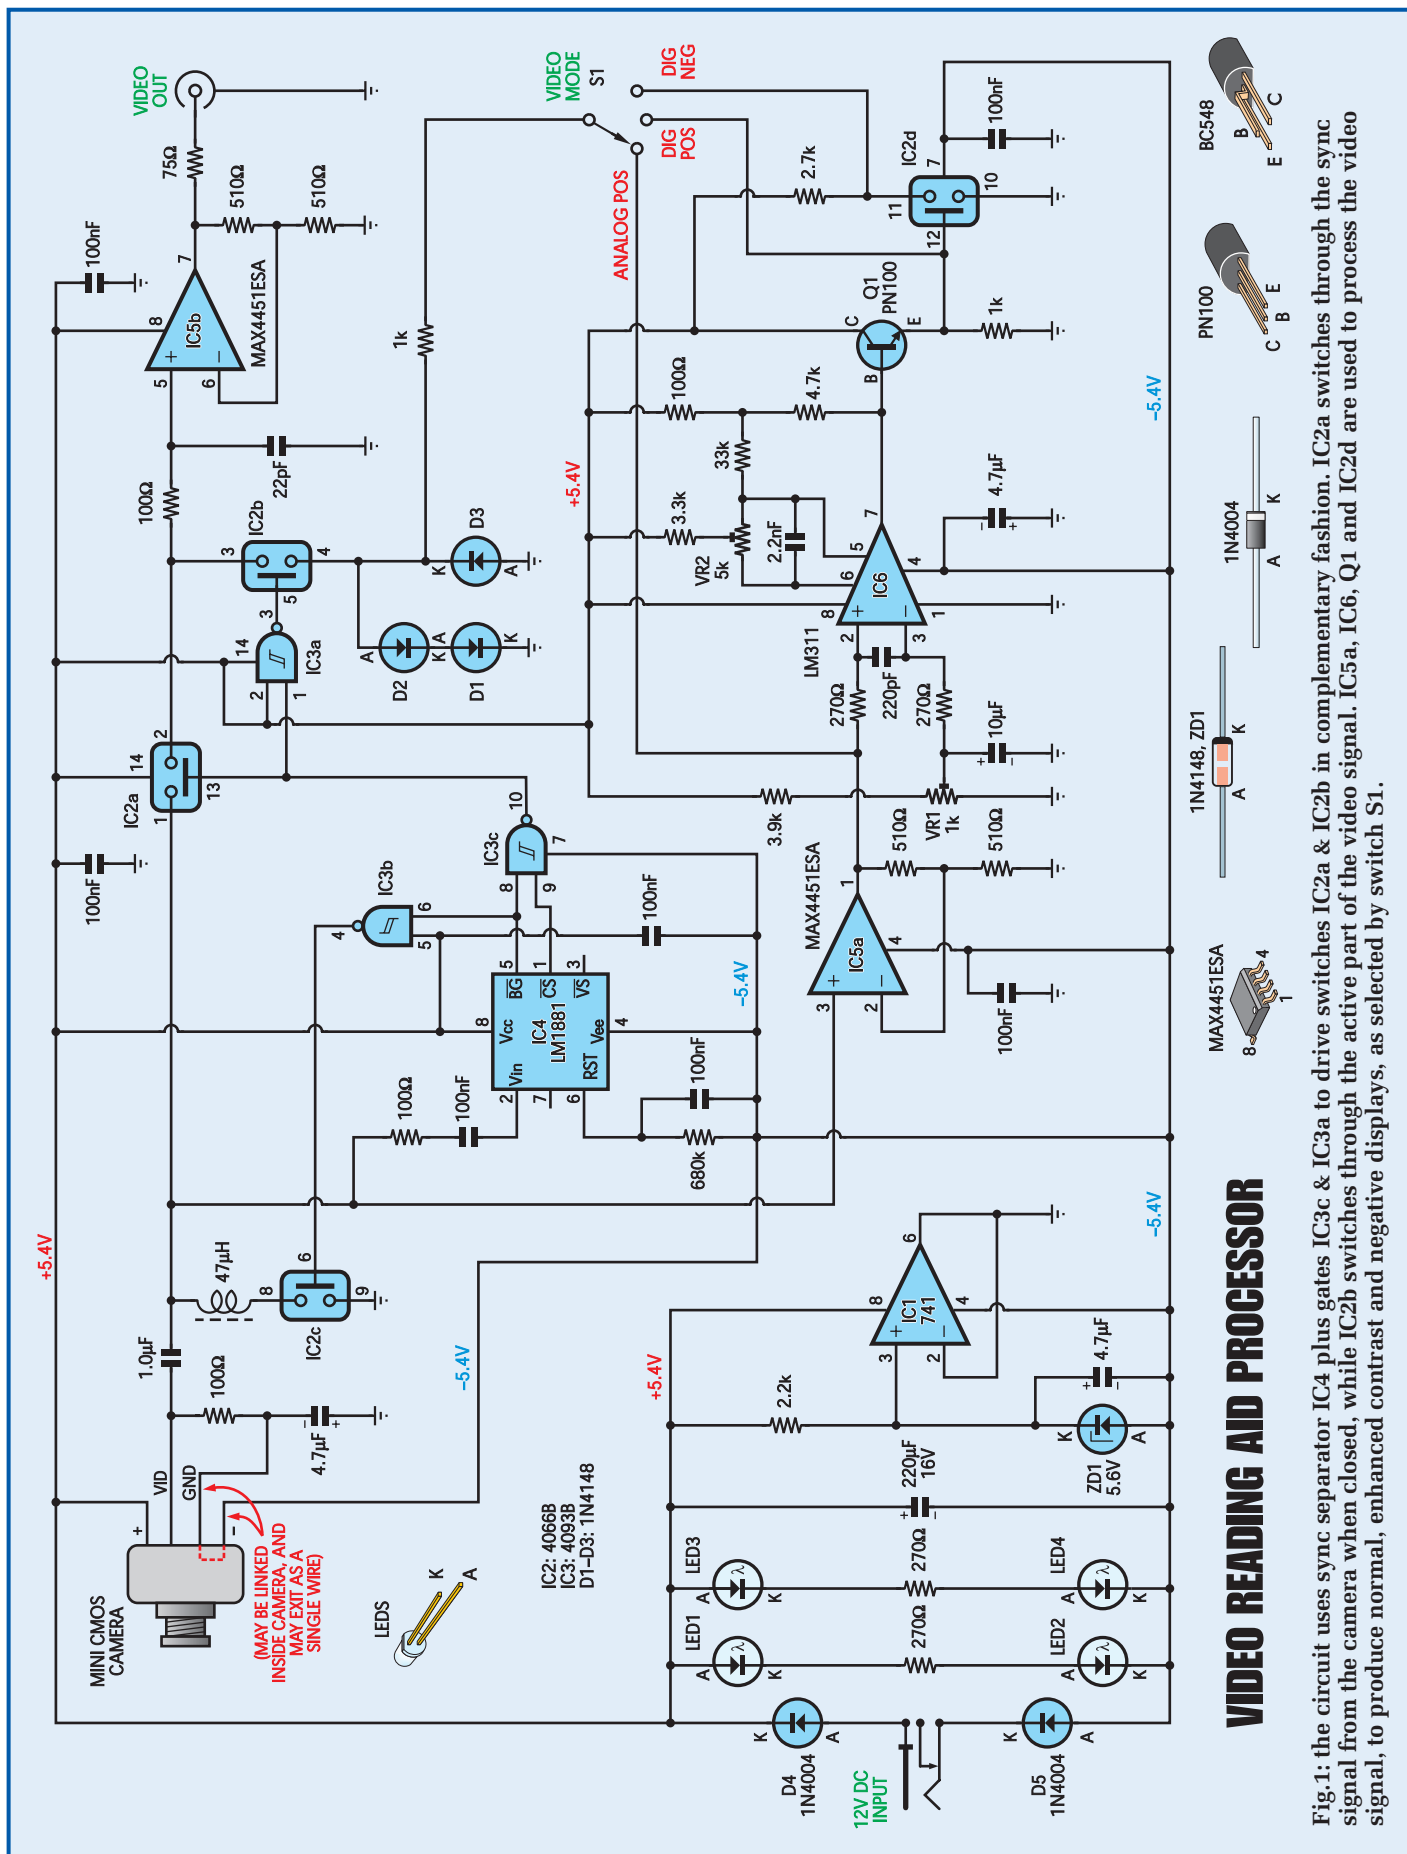
From there, the buffered signal is fed to the second position of selector switch S1, to become the hard limited or ‘Digital Positive’ video signal. This signal is also fed to the control gate (pin 12) of IC2d, used here as an analogue inverter. The inverted video signal appears at pin 11 and is fed to the third position of S1, to become the ‘Digital Negative’ video signal.

Limiting circuit

The processed video signal selected by switch S1 is first fed through a

simple diode limiting circuit involving diodes D1-D3 and a 1k Ω series resistor. Diode D3 ensures that the negative excursions of the signal (ie, its black level) are clamped at 0.6V below ground, while D1 and D2 ensure that the positive excursions (ie, peak white level) are clamped at 1.2V above ground. The processed video fed to video switch IC2b is thus limited to a fairly normal voltage range, so it shouldn’t cause any overload problems, either in the video output buffer stage (IC5b) or in the TV set.

The recombined sync and video signals from switches IC2a and IC2b are fed to pin 5 of IC5b via a simple low-pass filter comprising a series 100 Ω resistor and a 22pF capacitor. This removes any switching transients. The signals are then passed



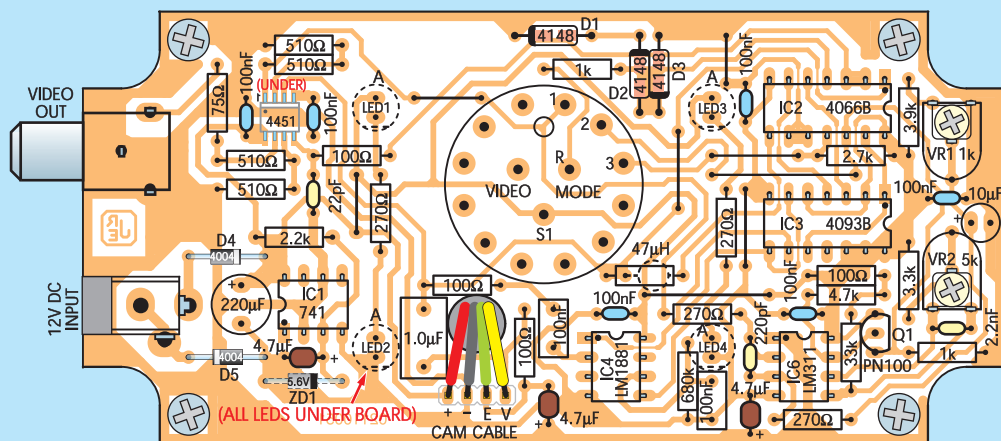


Fig.2: follow this assembly diagram to install the parts on the PC board, taking care to ensure correct component polarity. The four high-brightness LEDs and the MAX4451ESA device are installed on the copper side of the board (see Fig.3).



This is the fully-assembled PC board, mounted on 25mm tapped spacers. Note how the high-brightness LEDs hang down from the underside.

full lead length, so that the body of each LED protrudes through a matching 5.5mm hole in the box.

Fig.2 shows the parts layout on the PC board. Begin the assembly by fitting the 12V DC input and video output connectors, then install the wire links.

Next, fit the 4-pin SIL header which is used to terminate the leads from the CMOS camera. This goes just below the 8mm hole that the camera leads feed through. That done, you can begin fitting the passive components, starting with the resistors and RF choke and following these with the two trimpots, the smaller capacitors and finally the polarised tantalum and electrolytic capacitors.

Follow these with diodes D1 to D5, making sure you fit each one the correct way around as shown in Fig.2.

Also, make sure you use the larger power diodes for D4 and D5 and the smaller glass signal diodes for D1-D3. Zener diode ZD1 can then go in, again taking care with its polarity.

At this stage, it's a good idea to fit rotary switch S1. To do this, first cut its shaft to about 8mm long and carefully file off any burrs. That done, it can be mounted on the board with its indexing spigot at the 12 o'clock position, as shown on the overlay diagram. Push it all the way down onto the board before soldering its pins.

The next step is to fit IC1, IC4, IC6, IC3 and IC2, in that order. Note that the last two of these devices are CMOS ICs, so be sure to take the usual precautions to avoid subjecting them to electrostatic damage – ie, don't touch their pins, make sure the tip of your

soldering iron is earthed and solder their supply pins (pins 7 and 14) first. It's also a good idea to 'discharge' yourself by touching an earthed metal object before handling these devices or, better still, wear an earthed wrist strap.

The board 'topside' assembly can now be completed by fitting transistor Q1. Be sure to orient it as shown, then flip the board over so that you can fit IC5 – see Fig.3.

This IC is in an SOIC-8 surface mount package which measures only about 5mm square and has a pin spacing of just 1.25mm. It is just large enough to be soldered in place by hand, provided you take your time and work carefully.

This job requires a soldering iron with a very fine tapered bit, which is

also well tinned and clean. You should use fine gauge (ie, 0.8mm) resin-cored wire solder, to ensure there are no solder bridges between adjacent pins.

The best procedure is to hold the device in position using a wooden toothpick while you carefully solder one of its supply leads – either pin 4 or 8. This involves just touching the outer end of the device lead with the soldering iron and feeding on the solder, so that a tiny drop melts and bonds the lead to the pad underneath.

That done, you can quickly solder the other supply lead and then the rest of the leads. So the trick is to make one joint first, to hold the device in place while you solder all the other leads.

Doctoring the LEDs

Now for the LEDs. These are left until last because, as mentioned earlier, they first have to be ‘doctored’.

As supplied, the rounded end of each LED’s clear body produces a fairly narrow semi-focused axial beam. That’s fine for most applications but not this one, as this would produce very uneven lighting below the camera lens, with four bright spots separated by relatively dark regions.

The cure is simple – by sanding four small ‘flats’ on the end of each LED, its light output becomes much more diffused and this gives a more even illumination. Fig.4 shows the basic idea.

It’s quite easy to sand these flats by hand, because the LED bodies are moulded in a fairly soft ‘water clear’ plastic. A small piece of medium sand paper wrapped around a piece of flat wood will do the job quite nicely and you will only need seven or eight




Fig.3 (left): use fine-gauge solder and a fine-tipped soldering iron to install the SOIC device (IC5) on the underside of the PC board.




Fig.4: here’s how the four high-brightness LEDs are modified to diffuse the light.

passes to produce each flat at the correct angle (the exact angle isn’t critical, by the way).

Don’t try to polish the surfaces after sanding – just leave them with the after-sanding matt finish, as this gives better light diffusion.

After all four LEDs have been treated, you can fit them to the underside of the board. They must all be mounted at full lead length (ie, with the shorter cathode leads just entering their matching holes), so that they’ll later protrude through the holes in the bottom of the box when the board assembly is fitted.

Before actually installing the LEDs, it’s a good idea to fit 20mm lengths of 2mm sleeving over each lead, to prevent accidental shorts. You can use red sleeving for the anode leads and green or black sleeving for the cathode leads.

After the LEDs have been fitted, the board assembly can be completed

by attaching four M3 x 25mm tapped spacers (one at each corner), using 6mm long M3 machine screws.

Box preparation

The box needs to have a number of holes cut in the bottom and lefthand end of the base, plus two holes in the lid. The positions and sizes of these holes are shown in Fig.5.

Depending on the box, you may also have to cut away some of the plastic ribs moulded on the inside at one end, near the input/output connectors. This can be done using a sharp chisel.

Table 2: Capacitor Codes

Value	µF Code	IEC Code	EIA Code
100nF	0.1µF	100n	104
2.2nF	.0022µF	2n2	222
220pF	NA	220p	220
22pF	NA	22p	22

Table 1: Resistor Colour Codes

	No.	Value	4-Band Code (1%)	5-Band Code (1%)
□	1	680kΩ	blue grey yellow brown	blue grey black orange brown
□	1	33kΩ	orange orange orange brown	orange orange black red brown
□	1	4.7kΩ	yellow violet red brown	yellow violet black brown brown
□	1	3.9kΩ	orange white red brown	orange white black brown brown
□	1	3.3kΩ	orange orange red brown	orange orange black brown brown
□	1	2.7kΩ	red violet red brown	red violet black brown brown
□	1	2.2kΩ	red red red brown	red red black brown brown
□	2	1kΩ	brown black red brown	brown black black brown brown
□	4	510Ω	green brown brown brown	green brown black black brown
□	4	270Ω	red violet brown brown	red violet black black brown
□	4	100Ω	brown black brown brown	brown black black black brown
□	1	75Ω	violet green black brown	violet green black gold brown




Fig.5: here are the drilling details for the plastic case. It's best to make the larger holes by drilling small-diameter holes first and then carefully enlarging them to size using a tapered reamer.




Fig.6: this diagram shows how it all fits inside the case. Note that the lid should sit firmly on the switch indexing ring, to keep it in place when everything is screwed down.

Mounting the camera

Once you've drilled all the holes, the mini camera can be prepared for mounting. First, remove the two screws which attach it to its existing U-bracket, then cut the camera's output cable about 40mm from the body (or its mini connection plug). Remove about 15mm of the outer sleeving from the end, then separate the individual leads. In most cases, the positive power lead has red insulation, while the video lead has yellow insulation. The negative power lead usually has either black insulation or is in the form of a screening ground braid.

If the camera also has an audio output (many of them do), this is usually a wire with white insulation. This output is not used in this project.

After separating the various leads, strip about 5mm of insulation from the ends and tin the exposed wire ends, ready for connection to the 4-way header on the PC board. If your camera has a ground braid, this should be neatly twisted together, sleeved and tinned as well.

With a camera like the Swann unit, you also have to disable the inbuilt IR LEDs (originally intended for night illumination). That's done by removing the back of the case (it's usually attached by two tiny screws) and removing one of the LEDs – either by

cutting its leads with side-cutters or desoldering them from the internal PC board.

You don't have to worry about the others, because they're usually connected in a series string.


The camera can now be mounted inside the box using two small L-brackets, made from 1mm aluminium sheet – see Figs.6 and 7. The camera mounts between the brackets using the same two screws which held it in its original U-bracket.

It's a good idea to fit an M2.5 flat washer on each screw before passing it through the hole in the L-bracket and then fit an M2.5 star lockwasher on each screw before it enters its tapped hole in the side of the camera. This arrangement keeps the camera firmly vertical when both screws are tightened.

The camera mounting brackets are then attached to the box using M3 x 10mm machine screws, nuts and lockwashers.

Final assembly


Once the camera is mounted, the PC board (with its mounting spacers) can be lowered into position. Feed the camera cable through its board access hole as you go and make sure the four LEDs all pass through their respective holes in the box base. The board assembly can then be secured from



underneath using M3 x 6mm screws into the tapped spacers.

Finally, connect the camera cable leads to their respective header pins on the PC board. The positive power lead (red) connects to the leftmost pin, nearest the 1µF MKT capacitor, while the video wire (yellow) connects to the rightmost pin. If present, the audio wire (white) is left disconnected – just tape it up so it can't make contact with anything.

If there's a negative power wire (black) separate from the ground braid, solder this to the second pin from the left and connect the ground braid to the remaining pin – ie, the third pin from the left. Alternatively, if there's no separate negative power wire, simply connect the ground braid to BOTH of the centre pins.



Above: this close-up view shows the mounting details for IC5.

Below: the plastic skirt has a clearance hole for the camera lens and is attached to the base of the case using double-sided adhesive strips.

Above: the CMOS camera is attached to the base of the case and its leads fed up through a small hole in the PC board.

The only other possibility is that your camera may have just a black negative wire and no ground braid. In this case, connect the black wire to both centre pins instead.

Switch indexing

Before testing the Video Reading Aid, you have to set the rotary switch so that it has only three positions and not four.

To check this, fit its knob temporarily to the spindle and try turning it to see how many positions are available. If there's only three, you can relax. But if there are four, the switch will need to be reset.

To do this, first turn the switch anticlockwise to its end position and then remove the knob. That done, unscrew the mounting nut, and remove both it and the star lockwasher underneath. This will reveal the indexing stop washer, which you then have to prise up using a small screwdriver. The underside of this washer has a small spigot, which sits in one of the matching slots in the

switch body.

If you look closely you'll see that there are a series of numbers moulded into the switch body, between the slots. The idea is to find the slot between the numbers '3' and '4' and refit the indexing washer with its spigot in that slot. Check that the switch now has only three positions, then refit the star lockwasher and nut.

Fitting the plastic skirt

The plastic skirt fitted to the unit is actually an upside-down food container. The recommended unit measures 130 x 105 x 60mm deep and has an

indent in the centre of its base which provides clearance for the LEDs. The unit is also curved towards the sides, which means that it naturally clears the four corner mounting screws that go into the spacers.

Attaching it is hardly rocket science – just cut a hole in the centre to clear the camera lens, attach some double-sided tape to its base and attach it to the bottom of the box.

If you use a different food container from the one we used, then you may have to also drill holes to clear the LEDs and the mounting screws.

Testing

Now for the smoke test! First, set the rotary switch to fully anticlockwise (Medium Contrast), set trimpot VR1 to fully anticlockwise and set VR2 to its mid-range position. That done, connect the Reading Aid's video output cable to the video input of a TV set and apply power.

Note: you must use a 12V regulated plugpack or 12V battery. Do not use an unregulated plugpack, otherwise you'll damage the camera and IC5.

If all is well, you should see a bluish-white glow from the illumination LEDs underneath the Reading Aid box. Now place the unit on some printed material. The image will probably be quite blurry initially – just adjust the lens until you get the correct focus by rotating it clockwise or anticlockwise. This will have to be done by trial and error, since the plastic skirt is in the way when the unit is resting on a surface but it shouldn't take long to get it just right.

You may also have to adjust the brightness and contrast controls on the TV to get a good image.

If there's no image or none of the LEDs is alight, you've probably got the power supply the wrong way around. No damage will result from this – just reverse the connections and all should be OK. However, if the image does appear but only two of the LEDs are alight, the odds are that you've connected at least one of the LEDs around the wrong way.

Clarity and contrast

If all LEDs are alight and you have a clear image on the TV, turn the rotary switch to its centre position. The image will probably go very dark but if you turn trimpot VR1 slowly clockwise with a small screwdriver, it should gradually turn into a very 'contrasty' but still clear black-and-white image. The correct setting for VR1 will be quite obvious – just set it for maximum clarity and best contrast.

If you can't achieve this by adjusting VR1 alone, you may also need to adjust VR2 slightly one way or the other.

Once the correct settings have been found, try switching S1 to the third position (fully clockwise). The image should change into a high contrast negative, with black type on a white background turning into white type on a black background, which many




Fig.8: check your PC board for defects by comparing it with this full-size etching pattern before installing any of the parts.

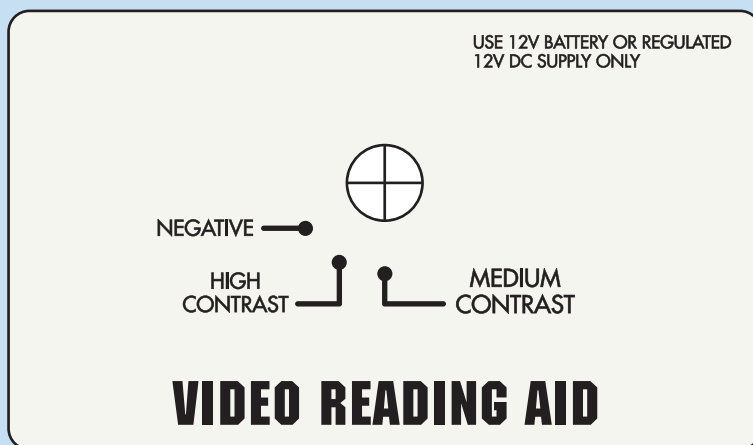


Fig.9: this is the full-size artwork for the front panel. It goes on the lid and can be protected using wide strips of clear adhesive tape.

people with visual impairment find easier to read.

Final assembly

Assuming it all checks out, disconnect the power supply and remove the knob, mounting nut and star lockwasher from the rotary switch. The box lid can then be slipped into position over the switch shaft and should rest on the top of the box, with the switch locating spigot passing up through the small hole that's located just behind the main spindle hole – see Fig.5.

All that remains now is to fit the four lid fastening screws and then refit the star lockwasher and nut to the switch ferrule. Your Video Reading Aid is now ready for use.

EPE

Image Washed Out?

Depending on the high-brightness LEDs supplied and/or the amount of ambient light at the reading location, you might find that the on-screen image is washed out (ie, over-bright).

In that case, try throttling back the LED brightness by increasing their series 270Ω resistors to around 680Ω. Alternatively, if you have plenty of ambient light, you may get a better result if the LEDs are taped over (or the unit modified so that they can be switched out of circuit).

Be prepared to experiment to get a good picture if necessary.

Reproduced by arrangement with
SILICON CHIP magazine 2007.
www.siliconchip.com.au

Using MPLAB

How to use MPLAB when writing the source code for your PIC projects Part Two – Initial Stages of Program Writing, by Mike Hibbett

LAST month we introduced the concepts behind MPLAB, and by now no doubt you will have had a tinker with the features.

During these early days it's possible to accidentally delete important files or perhaps corrupt the installation, requiring re-installation of the software. Or perhaps you want to upgrade to a newer version of MPLAB. Irrespective of the reason, it's probably sensible to explain how to re-install the application.

MPLab re-installation and upgrade

Make sure the MPLAB program is not running, then run the 'Control Panel' application from the Windows Start menu. Double click on the 'Add or Remove Programs' icon and wait for the list to be populated. Scroll down to the entry for 'MPLAB Tools' and click on 'Change/Remove'. At the next dialog click on 'Remove', followed by 'Next'.

Once the application has been removed, re-install by following the instructions from last month. Updates to MPLAB are all provided in a simple, single .exe installation file, so installing new versions will be just as simple.

Since last month's article, MPLAB has been upgraded from v7.52 to 7.60, so if you have high speed access to the Internet then now would be a good time to download the latest version (via www.microchip.com). Microchip follow a standard with version numbers on software releases; versions ending in '0' are major releases, others are minor. Version 7.60 should be Microchip's main release for a while. Version 7.52 is supplied on last month's cover-mounted disk, so we will stick with this version for now.

First simple example

The hardest part of any design is the 'blank sheet' at the beginning, so we begin gently now by walking through a simple example, highlighting the main features of the IDE. As mentioned in last month's article, we will start with an example of non-relocatable code development because this is simple to do and the most familiar to anyone who has used assemblers such as TK3 or MPASM. We will go into the more powerful relocatable development process in a later article.

So let's get started. Start the MPLAB program. The application window will appear with two child windows labelled 'Untitled Workspace' and 'Output'. The first thing we must do is tell MPLAB what type of processor we are programming for. MPLAB will use this selection to identify what kind of instruction set is permitted in your code, and what kinds of tools are supported.

For this tutorial we are going to use the PIC16F917. Don't worry if you are not

familiar with this part; the tutorial will be straightforward and the choice of processor will be useful later on when we demonstrate some of the advanced debugging features.

To select the device, click on 'Configure' from the main menu followed by 'Select Device...'. Choose 'All' from the 'Device Family' drop down list, then select the PIC16F917 from the 'Device' list. The red and green buttons below show the features that are available for your chosen processor. Green means supported, amber partially supported, and red not. Some features such as hardware debuggers and programmers are not supported yet for all processor types, but Microchip are working on it – hence the frequent software releases.

Close the dialog to return to the main window. Now click 'File' from the main menu followed by 'New'. A new window will appear titled 'Untitled'. This is a source file editor window, into which you can start entering your code. Click on the window to select it, and enter the program instructions from Fig.1.

The lines beginning with a semicolon are *comment lines* – they are ignored by the assembler. It's good practice to type them in, since they help explain the code. Enter the code exactly as shown – there are a few intentional bugs in it, so if you spot them, don't correct them yet!

```
#include P16F917.inc

; Select bank0 ram
bcf     STATUS,RP0
bcf     STATUS,RP1

; Specify the output values
; on PORTD
movlw   0x23
movwf   PORTD

; Select bank1 ram
bs      STATUS,RP0

; Make PORTD all outputs
clrf    TRISB

; select bank0 ram
btfss   STATUS,RP0

; go into a continuous loop
loop:   goto    loop

END
```

Fig.1. Setting break points

The MPLAB editor is quite advanced and has all the features you would expect from a source code editor (with the exception of a spell checker, unfortunately!). It may take time to get the hang of how to do the more advanced options (these are found in the 'Edit' menu option on the main menu) but with time it's easy to adjust.

Once you have entered the program, click on 'File' followed by 'Save As'. Use the 'Save As' dialog to save the file into a

convenient location, and give the file a '.asm' extension. Immediately, the edit window will change: the font changes and colours appear. The editor is 'context sensitive', meaning that it recognises when words are being used as instructions, comments or values and colours them differently to make them stand out.

Text formats

Notice how the instructions and comments are indented in the editor, but the label for the **goto** instruction ('loop') starts in column 1. The Microchip assembler assumes that any text which starts in the lefthand column is a label, and will generate a warning if it finds an instruction instead. This rather strange rule helps to simplify the design of the assembler software, so we just have to abide by the rule. It helps to make the code listing more readable anyway so is not an inconvenience.

Notice also that the instructions have been entered in lower case, while the names of SFRs (Special Function Registers) are in upper case. With the names of variables, labels and SFRs, the Microchip assembler is case sensitive; you must use the same combination of upper-case and lower-case letters. So for example:

STATUS

and

Status

refer to different variables or labels. It's best to avoid using such similar names in your code to avoid confusion – programming is confusing enough without needing us to add to the complexity!

The names for the instructions are shown in lowercase. We have done this purely for stylistic reasons – if you want to type instructions in uppercase, go right ahead. The most important thing is to find your own style and stick to it consistently. This will make re-reading your code at a later date much easier. For excellent advice on developing good programming style and techniques, take a look at the book *Code Complete* published by Microsoft Press.

INCLUDE files

Back to the source listing you just entered. The first line:

```
#include P16F917.inc
```

is called an 'include directive'. It isn't a PIC instruction, but a special command that instructs the assembler program to include the text found in the file 'P16F917.inc' when assembling your file. The file is a type of 'header file', always included at the top of a source file, that contains definitions of the various SFR and

bit names for the processor. The MPLAB program is supplied with a header file for each microcontroller supported. You do not have to specify the path to the file; the assembler knows where to look for it. If you would like to see the contents of the file for yourself, it is normally found in the following location:

c:\program files\Microchip\MPASM suite

There will be one for each processor type (so there are lots of them). You don't have to include this file into your source file; you could if you wish define the names for SFRs yourself. It is, however, a very good idea to do so. These files define all the registers in each microcontroller in a consistent way, and save you having to type them in yourself. The contents of the file will not add any code or data usage into your software, so just go ahead and include the file as a matter of course.

Processor type

Including the file **P16F917.inc** does not tell the assembler program what type of processor you are using; you did that in the IDE by clicking on the Configure/Select Device menu option. Therefore, when you close MPLAB down the information about the processor type will not be saved in your source file. MPLAB can create a special file, called a 'workspace', that remembers all the IDE settings, including what files you have been looking at and what the processor type is. To create this file click on 'File' followed by 'Save Workspace'.

You are prompted to name the workspace file and specify where the file should be located. Give it a name of **test.mcw** and save it in the same directory as your **test.asm** file. To prove that everything has worked, close the MPLAB program (answering 'Yes' to the question about saving the workspace). You can now click on the file **test.mcw** to start MPLAB restoring your files and settings exactly as you left them. That's it!

Program assembly

We are now ready to build (assemble) your first program. First, click on the **test.asm** window, and then from the main menu bar of the IDE select 'Project' and then 'Quickbuild'. The Output displays the progress of the various programs involved in creating your program. The final line of text written to this window should say:

'BUILD FAILED'

If you get the message 'Please put the .asm file that you would like to assemble into focus and try again', go back and click on the window containing your source file and try again. The Quickbuild option is designed for simple projects and does not really 'know' about the files in your project. We will cover the more advanced build options that do understand all your files in a later article.

Once the assembly stage has finished, take a look at the output window. Obviously, the program has some problems with it. Looking at the nine or so lines in the output window the one starting with

'Error' points to the serious problem. Double-clicking on this line of text will bring up the editor window and place the cursor on the line in question. This is the integration of the assembler and editor components of the MPLAB coming into play.

The error should be clear – we have miss-spelt the instruction. Change 'bs' to 'bsf' and re-build the program again. This time you should be greeted with the message:

'BUILD SUCCEEDED'

Warning messages

Take a look at the output window again. Although the error message has gone, a warning message is shown, which starts as:

Message[302]

Messages such as this (type 302) indicate possible problems with your program, but ones which are not serious enough to stop the assembler producing a **.hex** file. You should pay attention to all these messages; they will help point out the most obvious and trivial faults in your code.

Message[302] is a special case, and appears frequently. It's a completely benign message reminding you that a register you are accessing is not in **BANK0**. The assembler is not smart enough to know whether you have set the **BANK Select** bits appropriately, and so prints this message as a way to jog your memory. Many people consider this message to be an annoyance, so it's fortunate that it can be disabled. To do this, add the following line at the top of your source file:

errorlevel -302

Now when you rebuild the program the output window will display no errors and no warning messages.

The various different warning messages are described inside the help system of MPLAB. To view them, activate the help system by clicking on 'Help' followed by 'Topics' and then 'MPLAB Assembler'. Within the Index window pane, type in the word 'Messages' and click on the sub-heading 'Assembler'. You can also find the various error messages that can be displayed by typing in 'Errors'.

The explanation of 'errorlevel' can also be found here. Although it is possible to suppress the display of any error message, it is generally accepted that message 302 is the only one that can be safely ignored.

The fact that the assembler has successfully assembled the program is, of course, no guarantee that the program will work. The assembler has found no *syntax* errors; there could, of course, be plenty of *logic* errors. And this is indeed the case for our simple program.

Generated files

Before attempting to test the new program, let's take a quick look at the files generated by the assembly process. Four new files have been created, all with the same name as the source file – 'test' in this case, but with different file extensions. Two of them are of no interest to us at the moment (**.err** and **.cod**) but the **.hex** file

contains the raw program instructions in a format recognised by the programmer.

The last file, **test.lst**, is a text file that contains a listing of the program instructions alongside your original assembly code. At the bottom of the file is a memory usage map, which tells you how much of your flash memory has been used up – 8 out of 8184 words in this case. This is a very good indicator of how much space you have available as your program increases in size.

The list file allows you to see exactly where each part of your program has been located into the memory map. This can be very useful if you are coding 'lookup tables' or need to know in which pages of memory routines are located.

Simulation

The **.hex** file can now be programmed into a chip, either using your own programming utility or through MPLAB with a supported programmer unit. But this is jumping ahead somewhat; we do not really know yet whether the code actually works. This is where the simulator comes into play.

The simulator contains a model of what your processor comprises (CPU, memory, registers, peripheral features etc) which gets updated each time a line of your software is executed. Even on a very modern PC this process is very slow, taking several milliseconds to execute each one.

It is possible, however, to stop the execution at any point and see the full state of the processor. External signals on interfaces, like serial ports and I/O pins can be simulated, allowing a large amount of functionality to be tested without even wiring up a real processor circuit. It's not a complete substitute for the real thing – you cannot connect up your real external signals – but it can assist with finding those initial design errors quickly and easily. Let's see how it helps to find ours!

From the IDE main menu bar, click on 'Debugger', then 'Select Tool'. Click on 'MPLAB SIM'. Back on the main menu bar click on 'Debugger' again, then 'Run'.

After a few seconds, a dialog like that in Fig.2 appears. This is the classic error of not remembering to set the Config bits properly – the Watchdog Timer has been left enabled. Click on 'Yes' to close the dialog, and then click on 'Configure' followed by 'Configuration Bits'. Uncheck the flag 'Configuration bits set in Code', and then change the Watchdog Timer from On to Off. Close the dialog.

Rebuild the assembly program (to allow the Config bit changes to be incorporated into the program) and then re-start the debugger as described above. The warning message no longer pops up – that's the first problem solved – but the debugger appears to have stopped abruptly (the MPLAB window should continue to display the word 'running...' in the lower left of the display.) Something isn't quite right.

First, we need to put the debugger back into its initial state. Click on 'Debugger' followed by 'Reset' and then 'Processor Reset'. We now have several options for how we run the program; step, animate or run. Run, as we have already seen, simply runs the application as fast as it can, without giving any feedback. To find out what




Fig.2. Watchdog warning

is going on, you have to select the 'Halt' function in the debugger, which stops the code execution and displays the status of the processor.

The more useful option is 'Step Into', which causes the simulator to execute just one line of assembly each time the option is selected. 'Step Over' is very similar, but will not recurse into subroutine calls. The subroutine calls will be executed – they are just run at high speed. This is helpful when you are trying to locate errors in high level code that calls several routines known to be working ok. Typically, a debugging session will use both methods – you skip over known working subroutine calls, and step into the untested ones.

The last option for running a program in the debugger is 'Animate'. This option simply steps through the code slowly, a line at a time. It runs slow enough so that you can see the path taken through the program. To regain control, you click on the 'Halt' option, and can then step, run or reset the processor as you see fit.

Breakpoints

Under normal debugging sessions the code that you want to step through is buried deep inside the program, or perhaps lies in a section of code that is accessed infrequently. In cases like these, *breakpoints* are the solution. A breakpoint is simply a marker placed against a line of code that instructs the simulator to halt should it reach that line. Using one or more breakpoints you can start your software running at high speed, and then single step once the breakpoint has been reached.

Setting a breakpoint is simple – just double click on a line of code in the editor. A red 'B' icon will appear on the left hand side of the editor window, as shown in Fig.3. (The extra informational windows shown can be opened by clicking on the 'View' menu option.)

Let's complete this month's article by debugging the application to find the problem. Add the breakpoint as shown in Fig.3, and then click on 'Debugger', followed by

'Run'. The simulator runs briefly and then stops, with the green arrow positioned on the breakpoint line. This arrow indicates the next instruction that will run. Now, press the 'F8' button to step over each line. The arrow moves forward executing each instruction, until the line:

btfss STATUS,RP0

is executed – the arrow disappears! That's odd, the arrow should have moved to the next line. Something is wrong here, on this line. Ah ha! The comment on the line above gives the intent, but clearly we have typed in the wrong instruction. It should say:

bcf STATUS,RP0

Make that change to the program, re-assemble it and run the program in the debugger again. This time, the program enters the loop, as we would expect.

There is another problem with this program, and one that is not so simple to see in a simulator. It's a problem that becomes very evident when we run it on the target hardware. Debugging target hardware can be very difficult, but fortunately there are low cost solutions – tools such as the PicKit2 hardware debugger, which make solving problems on real hardware a breeze.

Next month we will demonstrate how such tools can be used, and solve the final bug. We will also look at using relocatable code generation for more complex programming problems

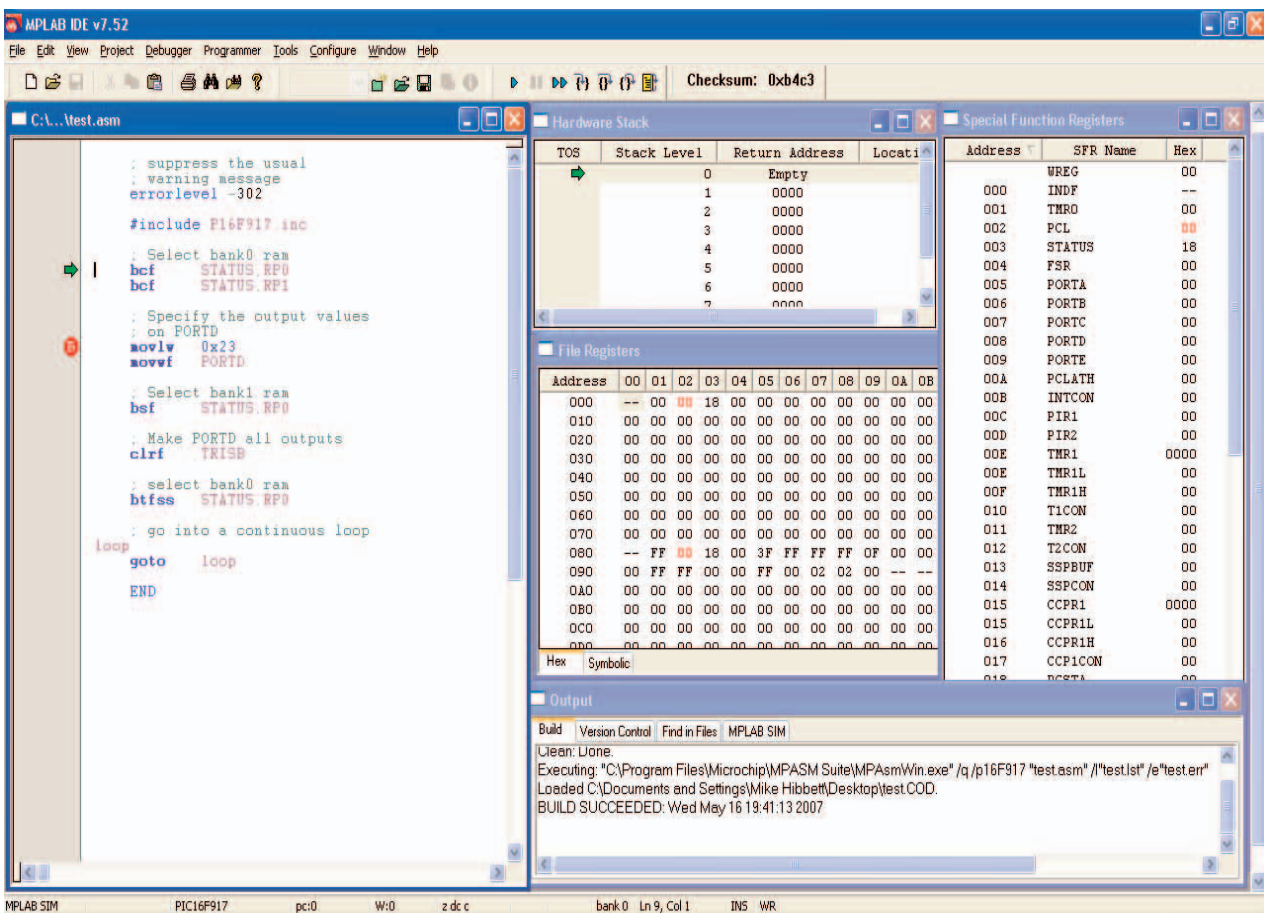


Fig.3. MPLAB source code listing screen

Circuit Surgery

Ian Bell

Linear Voltage Regulators and Capacitors – Part One

RECENTLY *Techno* posted the following question relating to linear voltage regulators on the *EPE Chat Zone* forum (via www.epemag.co.uk):

“Hi guys, I have noticed on voltage regulator circuits as used in power supplies that there are two non-electrolytic capacitors with a value of 100n wired between the input of the regulator and common ground rail and the other is wired between the output and common ground rail.

I am assuming the capacitors are used to smooth out excessive ripple from the supply. I have built power supplies in the past with regulators. I have not used the capacitor configuration in my designs and my power supplies worked well. I see the capacitor configuration is mainly used around the 78 series regulators. What is the purpose of the two capacitors?”

Stability factors

The issue of capacitances connected to regulators is more complex than a simple matter of smoothing. Regulators are feedback control systems which try to maintain the desired output voltage despite changes in current demand from the load. Feedback systems have the potential of becoming unstable, that is, they may oscillate, if the wrong conditions occur. For some types of regulator the choice of output capacitor is important in ensuring the stability of the circuit. The full characteristics of the capacitor (not just its basic capacitance value) may be critical to the stability of the circuit.

Some of the earliest linear regulators, such as the 7805 and 7815, were very stable under a wide range of conditions and did not require external capacitors. However, these regulators suffered from relatively high power dissipation due to the large voltage drop (typically a couple of volts) required for regulators to operate. To overcome this, different regulator circuit configurations were developed, which dropped only a few hundred millivolts. These are known as *low dropout* or *LDO* regulators.

Unlike the earlier circuits, LDO regulators require external capacitors to ensure stability of their feedback control loop, and for many chips the choice of the correct capacitor is critical to ensuring stability.

To understand what is happening with LDO regulator stability you need to understand a bit about feedback theory in

general. So, we will put the regulators to one side for a while and look at feedback and stability in the context of op amp circuits. Next month we will return to the regulators.




Fig. 1. General form of an amplifier with negative feedback

Feedback theory

The voltage gain of a typical op amp is very high (often a million or more), but op amps are usually used with *negative feedback* in various amplifier configurations, in which the amplifier circuit has much lower gain than the op amp. Negative feedback involves taking a fraction, β , of the output V_{out} , that is βV_{out} (where β is a value between 0 and 1) and subtracting it from the input signal V_{in} to give $(V_{in} - \beta V_{out})$ as the actual signal which is amplified by the op amp (see Fig. 1).

The schematic in Fig. 1 is very abstract and blocks do not necessarily directly represent real components. Block A is an amplifier with gain A. Block β simply outputs a fraction β of its input signal. The block drawn as a circle with a + sign adds its input signals to form its output. The minus sign by the input carrying the βV_{out} signal shows that this signal is subtracted rather than added.

The fraction β is called the *feedback factor*. If the gain of the op amp is A, then the output will be

$$V_{out} = A(V_{in} - \beta V_{out})$$

A is referred to as the *open loop gain*; it is the gain of the op amp itself, not that of the whole circuit. The gain of the whole circuit with feedback is called the *closed loop gain*, A_{CL} and is defined as

$$A_{CL} = V_{out} / V_{in}$$

where V_{out} and V_{in} are the circuit's output and input voltages.

Equation steps

We can rearrange the first equation. For the benefit of readers who are not confident messing about with equations we will show all the steps in detail.

$$V_{out} = A(V_{in} - \beta V_{out})$$

Divide both sides by A

$$V_{out} / A = V_{in} - \beta V_{out}$$

Add βV_{out} to both sides

$$V_{out} / A + \beta V_{out} = V_{in}$$

Collect terms in V_{out}

$$V_{out}(1/A + \beta) = V_{in}$$

Multiply β by A / A (=1) to facilitate writing $(1/A + \beta)$ as a single fraction

$$V_{out}(1/A + \beta A / A) = V_{in}$$

Take the common /A term outside the parentheses

$$V_{out}(1 + \beta A) / A = V_{in}$$

Divide both sides by V_{in} and then multiply both sides by $A/(1 + \beta A)$

$$V_{out} / V_{in} = A/(1 + \beta A)$$

So $A_{CL} = A/(1 + \beta A)$

This has given us an equation for the closed loop gain in terms of A and β . The quantity βA is called the *loop gain* and is important in determining stability, as we will see later.

Typical values

Now consider some typical values: if A is 1000000 and β is 0.1, βA is 100000 and $1 + \beta A$ is 100001. The 1 really does not make much difference. So, if A is very large, and the feedback fraction is not excessively small, we can ignore the 1 in $(1 + \beta A)$ and just write βA . Our equation for A_{CL} then becomes

$$A_{CL} = A/(\beta A)$$



But here we have A divided by A, that is the A's cancel out. So we get

$$A_{CL} = 1/\beta$$

This is a very profound result because it means that the gain of the circuit does not depend on the gain of the op amp, as long as the op amp has a very large gain. This is why the formulae for standard op amps are like R_2/R_1 or $1 + R_2/R_1$ – these formulae only contain resistor values – no parameters from the op amp at all!




Fig.2. Standard non-inverting op amp amplifier circuit

Loop gains

For closed loop gain A_{CL} to be independent of open loop gain A we need A_{CL} to be much smaller than A. This is usually not a problem. For example, if an op amp has a gain of 500000 and we require a circuit gain of 20, then we need $\beta = 0.05$, so $\beta A = 25000$, which is obviously much larger than 1 (our criteria for accepting the simplified formula $A_{CL} = 1/\beta$). The actual gain of the op amp if we use the full expression $A_{CL} = A / (1 + A\beta)$ will be 19.9992 instead of 20, a difference of 0.004% – compare this with typical resistor accuracy, for example 5%.

In Fig.2 is shown a standard non-inverting op amp. The op amp input voltages in this circuit are V_1 and V_2 , its output voltage is V_o and its gain is A. The circuit has a single input voltage V_{in} , an output voltage V_{out} and a gain of A_{CL} . For the op amp, $V_o = A(V_2 - V_1)$. This is because the op amp is a differential amplifier, amplifying the input difference ($V_2 - V_1$) by its gain A.

Gain answer

What is β for the circuit in Fig.2? Here R_1 and R_2 form a potential divider, which provides a portion of the output voltage at the op amp's negative input. The voltage at the negative input (V_2 in Fig.2) is given by the well-known potential divider formula

$$V_2 = R_1 V_{out} / (R_1 + R_2)$$

The voltage at V_1 is simply V_{in} , so for this circuit the op amp's output, which is given by $V_o = A(V_2 - V_1)$, can be written as

$$V_o = A(V_{in} - R_1 V_{out} / (R_1 + R_2))$$

which on comparison with our feedback formula above, $V_{out} = A(V_{in} - \beta V_{out})$, indicates that

$$\beta = R_1 / (R_1 + R_2).$$

This expression for β should not be surprising, as it is simply the proportion of the output provided by the potential divider. If our 'high op amp gain' assumption holds, we can write the circuit gain as $1/\beta$, which is $(R_1 + R_2) / R_1$ or $1 + R_2 / R_1$. Thus, as previously stated, the gain of the circuit is determined by R_1 and R_2 and is independent of the op amp's gain, so long as that is high, making circuit design of the amplifier very straightforward.

Response delays

The output of an amplifier does not respond infinitely quickly to changes at its input, so any signal fed back from the output to the input will be offset in time with respect to the original input. Consider a simple case in which there is a fixed delay from input to output of the amplifier, whatever the input signal does (things are usually more complicated than this).

Say, for example, this delay is $0.1\mu s$. If the input frequency is 100Hz, this time would be 0.001% of the signal's cycle time and could probably be considered insignificant. However, at 2.5MHz the $0.1\mu s$ delay is a quarter of the signal's cycle time of $0.4\mu s$. This would usually be expressed by saying that the amplifier has a phase shift of 90° at 2.5MHz (one complete cycle of the waveform is 360°). At 5MHz, $0.1\mu s$ is half the cycle time of the signal. This is a significant point because a phase shift of 180° is equivalent to multiplying the signal by -1 .

tend towards infinity. It is this infinite closed loop gain which results in instability. The condition for which $(1 + \beta A) = 0$ is $\beta A = -1$. More specifically, as indicated above, we get instability when the magnitude of βA is 1 (we write this as $\beta A = 1$) and the phase shift due to βA is $\pm 180^\circ$. Obviously, oscillation is undesirable in an amplifier circuit so we need an understanding of how it might occur in order to design to prevent it.

Avoiding instability

In most op amps, as the frequency is increased, the gain decreases and the phase shift moves towards $\pm 180^\circ$ from a smaller value. The circuit may not oscillate, however, because as frequency increases either:

- 1. As βA approaches 1 the phase shift is less than 180° . The difference between the phase at this point and 180° is the *phase margin*.
- 2. As the phase shift of βA approaches $\pm 180^\circ$, the magnitude of βA may be less than 1. This difference can be expressed as the *gain margin* (usually in dB).

Gain margin and phase margin are illustrated in Fig.3, which shows the variation of magnitude (gain) of A and (phase shift) of βA with signal frequency. Note that a gain of 1 is 0dB and that phase shift is negative, because the output lags behind the input signal in time.




Fig.3. Variation of magnitude (gain) of βA (dB) and phase shift of βA with signal frequency, illustrating gain margin and phase margin

Consider the total phase shift through the amplifier and feedback network as we increase the input signal frequency. Once this reaches 180° we have effectively inverted our feedback signal – what was negative feedback has become positive feedback. Positive feedback is what you need to make an oscillator, so our amplifier may become unstable.

For this instability to occur, the loop gain βA must be 1 or more. Recall the basic equation for closed loop gain with feedback β and open loop gain A

$$A_{CL} = A / (1 + \beta A)$$

Now, if the value of $(1 + \beta A)$ tends towards zero the closed loop gain will

The larger the feedback fraction β the more 'difficult' it is to fulfill the gain and phase margin stability criteria because the loop gain is higher. Thus, a circuit could, for example, be stable with $\beta = 0.5$ but not with $\beta = 1.0$. The decrease in gain with frequency from an op amp is not arbitrary, it is part of the design of the op amp.

Introducing circuitry to modify how gain and phase shift change with frequency in such a way as to ensure stability is known as *compensation*. The capacitors in the regulator circuits which prompted this discussion are used to compensate the regulator's feedback system to ensure stability.

Next month, we will look in more detail at how compensation is actually achieved in both op amps and regulators.

The Power of Mechatronics

Part Two – Controlling Motors by Darren Wenn

THIS article is the second in a series dedicated to looking at the science of mechatronics and how PIC microcontrollers can be simply used to provide increased functionality, higher performance and lower cost solutions in common applications.

The first question that then arises is what exactly is mechatronics? It can be put simply, as the application of modern electronics and software to conventional mechanical systems so as to obtain improvements in design, efficiency, functionality, size, reliability and cost.

Mechatronics technology

Mechatronics technology has been applied in a variety of areas from automotive applications through to home goods. In recent years readers will no doubt have noticed the explosion of ‘must-have’ gadgets – who could possibly survive without an electric toothbrush with integral LCD display? Even the simple light switch can be enhanced with the addition of a 6-pin PIC10F microcontroller, allowing it to monitor the current, provide dimming functions via a triac and even indicate when the light bulb has blown, although one suspects that the last point would be pretty easy to spot!

One very common aspect of mechatronics is it will often seek to replace electro-mechanical systems with a microcontroller enabled version driving the same actuators and motors but simplifying the mechanical aspects of the design. For example, an electric seat in an automobile could have the multipart lever and catch assembly replaced with a single motor and lead screw controlled by a PIC microcontroller.

So since it appears that driving actuators is an important aspect of mechatronics, this article is going to concentrate on various techniques that can be used to provide efficient and precise control of common electric motors using the Microchip PICDEM mechatronics demonstration board.

Talking back

Before we begin to look at how to control an electric motor using a PIC microcontroller, it is worth taking a small detour. Whenever we are developing any kind of system using a microcontroller it is useful to have the ability to provide some form of feedback to the developer on how it is operating. We could use a simple LED to

Listing 1

```
void main(void){
    unsigned long input;

    PIE1      = 0b00000000;
    INTCON    = 0b00000000;

    // set up the serial port for RS-232 comms so that we can get
    // feedback from the motor when it is running
    OSCCON    = 0b01110000; // change to 8MHz operation
    ANSEL     = 0b00000001; // make AN0 analog input
    TXSTA     = 0b00100100; // High speed, Async mode, TX enabled, 8-bit
    RCSTA     = 0b10010000; // RX enabled, Enable module
    SPBRG     = 51;         // 9600 baud
```

indicate the operating mode but its ability to signal complex data is somewhat limited.

Within the MPLAB development environment there is, of course, a highly capable In-Circuit Debugger (ICD) that connects to the mechatronics board via the MPLAB ICD 2. This is ideal for analysing the program code and stepping through the program while it is running on the target device, and in the unlikely event of a bug being present it provides a simple method for tracking it down and fixing the system.

However, when using fast moving dynamic systems such as motors, it is frequently necessary to provide real-time feedback without halting the motor. To this end we can use the EUSART (Enhanced Universal Synchronous Asynchronous Receiver Transmitter) to communicate over an RS232 link to an attached PC.

For our demonstration system we are going to use the PIC16F917 microcontroller which is provided with the mechatronics demonstration board. For compactness the software listings are written in C using the HI-TECH PICC-Lite compiler which is available from Microchip (www.microchip.com) for trial evaluation. Other C compilers would suit equally well and for those with time on their hands then the software could be written in assembler. In Listing 1 is the C code

for configuring the EUSART for 9600 Baud 8N1 communications.

Mechatronics and Motors

A fundamental aspect of mechatronics is the ability to control electric motors or actuators, so let's take a quick look at the physical make-up of a typical Brushed-DC motor. Unlike more modern alternatives, such as the Brushless-DC motor, the Brushed-DC (BDC) is the kind of motor often found in simple electronic devices and is familiar to most people. Its basic structure is shown in Fig.1.

The outside of the BDC motor, called the *stator*, consists of pairs of opposing permanent magnets that generate a stationary magnetic field. Inside the body is a rotating rotor, or armature. The rotor has one or more windings of wire wrapped around it. As current is fed through the windings it induces a magnetic field which




Fig.1. Simple Brushed DC motor

interacts with the permanent magnetic field of the stator. This interaction creates a torque which will cause the rotor to turn hence generating motion.

Of course, the problem is how to connect to the rotating windings and this is done by connecting them to a ring of metal cut in various places and called a *commutator*. As the commutator rotates, fixed brushes, usually made of carbon, rub against it and provide a way for electricity to pass into the rotor. As the motor turns, the windings are constantly energised in a fixed sequence causing continual motion of the rotor.

The main problem with BDC motors is related to this ring and the brushes. Both parts wear out over time due to the friction as they rub against each other. As the commutator changes windings a small spark is also generated which can prevent the motor being used in dangerous applications, and it emits a considerable amount of electromagnetic interference. Even with these problems the BDC motor is still the most common type of motor available for the hobbyist.

Driving around

So how do we go about connecting a BDC up to a PIC? It is usually not a good idea to connect a motor directly to a micro, as the high voltages and high currents used in motors are likely to damage it. It is normal to provide an external drive circuit that is capable of handling the higher power levels required by the motor.

In the case of the BDC motor it is only necessary to vary the current in the windings since the actual sequencing, or commutation, is handled by the brushes and metal ring inside the motor. This contrasts with some more modern motors, such as the Brushless-DC type, which require complex techniques to electronically perform the commutation as well as controlling the current.

One of the simplest drive arrangements is to connect the motor directly to V_{cc} and then connect the lower end to ground via a MOSFET. The low effective on-resistance of the FET ($R_{ds(on)}$) allows for large current flows with minimal power dissipation in the control electronics. The basic arrangement is shown in Fig.2.

Notice that if the FET is turned off when the motor is still turning it will start to act as a generator building up potentially high voltages that could cause damage to the FET. To prevent this, D1 acts as a free-wheeling diode to dissipate the excess energy.

With the simple low-side drive, a FET driver is frequently not required. However, the disadvantage is that the motor is permanently connected to the high voltage supply. Should the FET fail or a short occur, then the motor will be powered. Also, it may not be physically possible to isolate the negative (ground) supply connection from the motor, as in the case of many automotive motors where the casing forms the ground connection.

To overcome these problems, a high-side drive could be used. Also shown in Fig.2, the high-side drive has the FET located above the motor so that one terminal of the motor is permanently grounded. While this removes the grounding problem, the drive circuit may become more complex, since it is now necessary to convert the PIC output signal to a high voltage using a MOSFET driver in order to turn on the P-type FET.

When the PIC starts up, its I/O pins will initially go into a high impedance state. Depending on the drive circuit this high impedance state may be enough to turn the motor on, so to prevent this happening, resistor R2 biases the MOSFET to an off state. Resistor R1 is used to prevent excess current spikes damaging the micro.

The need for speed

We now have a very simple drive circuit capable of switching a BDC motor on and off from a PIC. However, this does not get us much further than a conventional switch, so now let us apply some mecha-

tronics and start to vary the speed of the motor using the PIC.

The speed of a BDC motor is proportional to the voltage applied to the motor and the torque generated is directly proportional to the current passing through the windings. To produce a controlled voltage in the motor we could simply regulate it using a series resistor or a transistor. However, these methods are not very micro

friendly and can dissipate large amounts of power. We can do a much better job by pulse-width modulating (PWM) the gate drive signal.

If we apply a PWM signal to the motor, the windings will act as a low-pass filter so that the motor sees a lower average voltage. If a motor runs at 2200rpm from a 5V supply and we drive the motor using a PWM signal with a 25% duty cycle, then the average voltage seen by the motor will be $5 \times 0.25 = 1.25V$ and will turn the motor at 550rpm.

While the duty cycle of the PWM will affect the apparent voltage and hence rpm, we also have to consider the frequency. If it is set too low then the motor may react slowly to changes and also have an annoying audible hum. Also, the upper limit of the frequency may be limited by the dynamics of the motor and the FET drive circuit.

At high switching frequencies the time taken for the FET to turn on and off becomes significant. It is during these times that the FET will dissipate most power (so called switching losses). A frequency in the range 4kHz to 20kHz is often chosen since this will reduce the amount of audio noise while keeping the switching losses in the FETs low.

Provided that the switching frequency is low, it is quite possible to use a PIC to toggle an output pin in software so as to achieve the desired PWM output. However, a simpler method is to use the CCP (Capture Compare/PWM) peripheral module provided on many of the device variants. This versatile peripheral can be used for things such as counting the duration of input pulses or setting outputs based upon internal timer values all without software intervention.

Of most use to us, though, is that it has a dedicated PWM output mode. The internal timer period register PR2 is loaded with a value that governs the PWM frequency according to:

$$PWM_{period} = (PR2 + 1) \times 4 \times T_{osc} \times T2_{prescale}$$

where T_{osc} is the PIC's oscillator period, which is 125ns when running at 8MHz. So with a value of $PR2 = 99$ the PWM will operate at 20kHz. The maximum allowable duty cycle is related to the frequency at which the PWM is operating, and is written into register CCPR2L. If the value in this register is higher than the timer period register, then the FET will be turned on permanently and the motor will run at full speed.

To allow the user to select the speed, we can set up the analogue-to-digital converter (ADC) to read the value of a potentiometer and use it to set the duty cycle register. Since the ADC reading is returned as a 10-bit value, we simply discard the lowest two bits and scale the result so that it has the same range as the timer period register. The code to configure the CCP module along with the ADC is shown in Listing 2.

Also shown is the main loop which simply reads the ADC, scales the result, outputs the new value to the serial port and updates the duty cycle register.

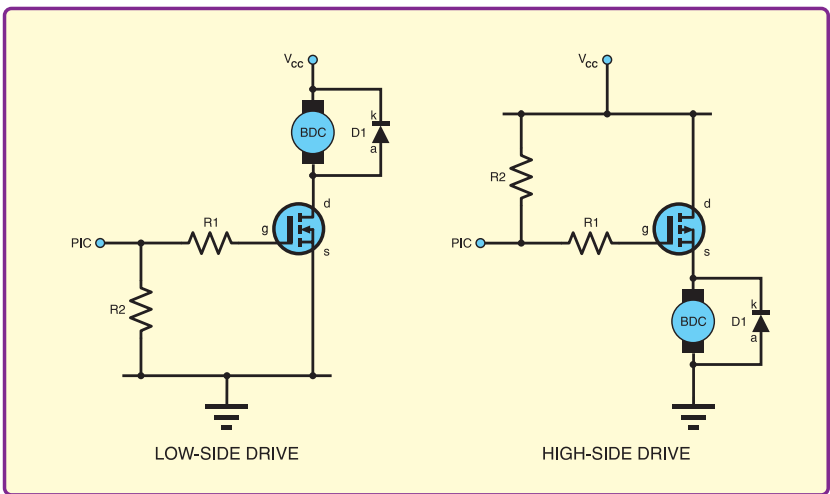


Fig.2. Brushed DC drive circuits

Listing 2

```
// set up the PWM output module for single direction speed control
PR2 = 99;           // 20 kHz PWM frequency
CCP2CON = 0b00001100; // single output PWM mode
CCPR2L = 48;        // set for initial 50% duty cycle (PR / 2)
TMR2ON = 1;         // turn on the timer and start the PWM
TRISD2 = 0;          // set pin as output

// set up to read POT1 using the ADC on Channel 0
ADCON0 = 0b00000001;

// main loop reads the ADC value and sets the duty cycle
while (1) {
    GODONE = 1;      // start an ADC Conversion
    while (GODONE);  // wait for it to finish
    input = ADRESH;   // read the result
    input *= 99;       // scale the result into the range 0-100
    input /= 255;
    printf("Speed is %d\r\n", input); // display on PC
    CCPR1L = input;
}
```

In order to run this code on the PICDEM mechatronics board, wire links should be made in the following places:

- J4:POT1 to J13:AN0 for the speed setting potentiometer
- J10:CCP1 to J1:N2 for the low-side drive FET
- J1:P1 to J10:V_{dd} to permanently turn on the high-side FET

These connections will allow the motor speed to be governed by the potentiometer setting. If you examine the circuit diagram of the mechatronics board, which is provided with it, then it can be seen that the drive FETs are all configured as high and low pairs. This arrangement is known as a *full-bridge*. We will consider this arrangement a little later on, but for now, the connections ensure that it will operate like a simple low-side drive.

Feedback

Now that we have built a motor controller, we could sit back and earn a nice profit. However, the motor controller as it stands is not ideal. The most important problem is that the control is very non-linear. If you examine the output of the program using a program such as HyperTerminal and watch the motor closely you will see that for a considerable portion of the potentiometer's rotation the motor does not turn.

This poor low speed performance is mainly caused by the friction of the motor. For a typical board this inactive region might extend up to 20% of the total available values. Ideally, we would like some way of kick-starting the motor to get it spinning at these low RPMs and then slow it down just enough to keep it running.

A second related problem is that if you load the motor by gripping the flywheel with your fingers you can hear the motor slow down, even though the speed command signal remains the same. We would like some method of applying extra torque

when the motor is loaded so that the speed is kept constant. In short we are going to need *feedback*.

For the BDC motor on the mechatronics board we can obtain feedback in a couple of ways. First, we can use the optical interrupter, which consists of an infra-red LED and a phototransistor which sees through the slots in the flywheel. As the motor rotates, this provides a pulse train directly proportional to the motor speed. We can compare this to the commanded speed and adjust the duty cycle to keep it constant.

A second method is slightly more complicated and involves measuring the *back electromagnetic flux* (BEMF) generated by the motor. If we configure the ADC to measure the voltage when one side of the motor is floating and the other is grounded, then the motor will be acting like a generator. In this case the voltage seen will be proportional to the speed of rotation and this can be used in the feedback calculations.

For this first part of the article we will configure the PIC to record the time between optical interrupter pulses. Once

scaled to a similar range to the potentiometer from our first example, the measurements can be used to generate an error signal. A proportion of the error signal is then added to the PWM duty cycle command signal and a new output is generated. If this program is run with the potentiometer set at a mid-value then the motor will run up to speed.

When a load is applied to the flywheel by gently gripping it, then the PIC will react to the motor slowing down by increasing the duty cycle of the PWM, and this should ultimately restore the RPM to its no-load speed.

You may notice some oscillation in the RPM and this is due to our very simple control algorithm. For readers interested in better performance then you would be well advised to look at any one of the vast number of articles on PID controllers that can be found on the web (many insomnia sufferers also find this literature useful!).

This simple algorithm does a reasonable job of keeping the motor speed constant, but the low speed performance is not brilliant. The software as shown implements a simple 'kick-start' scheme that detects when the motor has stalled at low speed and fires a short burst of high duty-cycle PWM into the motor to start it turning. The final version of this code is shown in listing 3.

Change of direction

Now that we can control the speed of the motor, and to a limited extent control its performance under load, the final thing that we are going to look at this month is how to control the direction of rotation of the motor.

Our simple low-side and high-side drivers worked quite well for speed control, but we need to be able to control the direction of current flow through the motors. It turns out that for a single rail supply there is only one practical circuit that will perform this task. This configuration of FETs and motor is known as an *H-Bridge* and can be seen in Fig.3.

When the motor is operating in a forward direction, FETs TR3 and TR2 are turned off, while TR1 and TR4 are PWM modulated. Current will then flow down




Fig.3. H-Bridge motor drive

Listing 3

```
volatile unsigned long actRPM;
volatile unsigned long tmrStall = 0;
volatile unsigned long prevTime = 0;
```

```
void main(void) {
    long cmdSpeed, IRPM, IErr;
    int cmdSignal;

    // RS232 and PWM module set up as before
    ...

    // set up CCP1 to time the duration of the dark pulse
    from the opto
    CCP1CON = 0b00000101; // capture every rising edge
    CCP1IE = 1;
    PEIE = 1; // enable CCP and
    peripheral interrupts
    OPTION = 0b00000111; // T0 prescale 256
    T0IE = 1; // enable T0
    T1CON = 0b00110001; // configure T1 for prescale 8
    GIE = 1; // enable interrupts
    cmdSignal = 0;

    // main loop calculates the duty cycle
    while (1) {
        GODONE = 1; // start an ADC Conversion
        while (GODONE); // wait for it to finish
        cmdSpeed = ADRESH;
        // scale the result into the range 0-100
        cmdSpeed = (cmdSpeed * 99) / 255;
        printf("%6d,", cmdSpeed);

        // convert counts into percentage of full scale RPM
        IRPM = 6010 / actRPM;
        IErr = cmdSpeed - IRPM;
        cmdSignal += IErr / 2;
    }
}
```

```
// limit the command signal
if (cmdSignal < 0) cmdSignal = 0;
if (cmdSignal > 100) cmdSignal = 100;

// handle stall conditions at slow speed
if ((tmrStall > 20) && (cmdSpeed > 0)) {
    tmrStall = 0;
    cmdSignal = 35;
}

printf("%6d, %6d, %6d\n", IRPM, IErr, cmdSignal);
CCPR2L = cmdSignal;
}

// read the capture port and store as a 10-bit result the time for
the last half revolution static void interrupt isr(void)
{
    unsigned long newTime;

    if (CCP1IF) {
        newTime = (CCPR1H << 2) | (CCPR2L >> 6);
        if (TMR1IF)
            actRPM = newTime + (1024 - prevTime);
        else
            actRPM = newTime - prevTime;
        prevTime = newTime;
        CCP1IF = 0;
        TMR1IF = 0;
        tmrStall = 0;
    }

    if (TOIF) {
        RB0 = !RB0;
        tmrStall++;
        TOIF = 0;
    }
}
```

through TR1, passing through the motor and on to ground via TR4. When the motor is operated in the reverse direction, TR1 and TR4 are turned off, and TR3 and TR2 are modulated. Depending on the particular system, the lower side active FET may be left switched on when operating in a given direction since modulating the upper FET is sufficient to control the motor.

There is another nice feature of the H-Bridge that allows rapid stopping of the BDC motor. If TR1 and TR3 are turned off, and TR2 and TR4 are turned on, then the motor will act as a generator, but it will appear as if it is connected to an infinite load, thus bringing the motor to a rapid halt.

The H-Bridge circuit is often found in commercial drives and by using high power FETs large motors can be controlled. A major problem can occur, though, if either of the side pairs of FETs are simultaneously turned on. It can be seen that switching TR1 and TR2 (or TR3 and TR4) will cause a direct short to

ground from V_{supply} which would ultimately destroy the FET.

We could apply some software logic to prevent this, but this additional work will degrade the performance of the controller. A better alternative is to use a microcontroller with dead-band compensation. This technique ensures that a delay is inserted when the upper FET turns off and the lower one turns on, and vice versa, so preventing the damaging 'shoot-through' as shown in Fig.4.

Whilst we have been using the PIC16F917 in our experiments so far, a PIC16F690 is also supplied in the mechatronics kit and can be inserted in the lower socket. This device features an Enhanced CCP (ECCP) port which is capable of driving an H-Bridge and also allows the dead-time to be programmed. For interested readers, Lab 9 on the CD accompanying the PICDEM mechatronics demo board shows how to set up and use the BDC in this bi-directional mode.

The finish line

This month we have seen how to drive a brushed DC motor. Whilst this is a simple electro-mechanical device, by controlling it from a PIC we have been able to vary its speed, provide relative position feedback and control its behaviour under load. All of these things would be harder to achieve using discrete circuitry and would certainly take longer than simply modifying a few lines of code. This is a practical example of mechatronics in action.

Now that we have laid some groundwork, in next month's article we are going to perform a little bit of 'hacking' and show how it is possible to fit a high performance 16-bit dsPIC DSC into the mechatronics board. This advanced controller will allow further experimentation and we will look at its motor control PWM module and show some digital signal processing in action.

Exclusive board offer

The Microchip PICDEM Mechatronics Development Board not only supports all of the projects featured in this series of articles but also includes nine example projects, each complete with source code.

To claim your exclusive *EPE* 20% discount on the Microchip PICDEM Mechatronics Development Board contact ACAL Semiconductors on Telephone: +44 (0)118 902 9702. Fax: +44 (0)118 902 9614. Email: sales@acalsemis.co.uk. Website: www.acalsemis.co.uk




Fig.4. Dead-band compensation

PRACTICALLY SPEAKING

Robert Penfold looks at the Techniques of Actually Doing It!

UNLESS you are keen on constructing valve radios and other projects from the distant past, at least one semiconductor will be used in every electronic project that you build. Semiconductor is a generic term that covers everything from a simple diode to the latest microprocessor and memory devices that contain the equivalent of many millions of components.

Modern electronic projects tend to be based on one complex integrated circuit (IC), which can be in the form of a micro-controller or a special chip designed specifically for one application.

Spoil for choice

The simpler semiconductors such as transistors and diodes are still used in modern electronic circuits, but to a lesser degree than in the past. It is probably fair to say that the simpler integrated circuits, such as operational amplifiers (op amps) and the basic logic types, are also used much less than previously. While the more basic semiconductors are less widely used, and the range of available types has contracted somewhat in recent years, they are still far from being extinct. Overall, the range of semiconductors on offer still remains vast.

Finding the right semiconductors in component catalogues can be problematic for beginners. Semiconductors are usually grouped into several broad categories in catalogues, with many of these being divided into several sub-categories.

You have to look for each device in the right category or sub-category in order to stand any chance of finding it. A full list of stocked devices is sometimes provided, but even here it is possible to miss the device that you are seeking.

Unless you know what you are doing, it can be slow and difficult to locate a device from a list having what could be a few thousands entries. Rather unhelpfully, some devices are produced by more than one manufacturer, with each producer using their own version of the type number. This makes it easy to overlook the right chip because it has a type number that is slightly different to the one that you seek.

Just a second

It is much easier to locate the right component if you understand the fundamentals of integrated circuit type numbering. No doubt there are exceptions, but practically all integrated circuits have type numbers that break down into three sections. The first of these usually consists of two or three letters that indicate the component's manufacturer. This is complicated by the fact that a manufacturer may use more than one set of letters. For example, linear devices might have a different prefix to logic types.

Many integrated circuits are second-sourced, which simply means that they are produced by more than one manufacturer.

Many industrial customers do not like being tied to a single source of supply, so some integrated circuits are manufactured under license by two or more additional manufacturers.

This makes the components more attractive to the buyers, but it gives rise to the problem mentioned previously. While the second-source components sometimes retain the original type number in full, often the prefix is changed to that of the second-source manufacturer.

What this means in practice is that you can often ignore the first part of the type number and concentrate on obtaining one where the rest of the type number is precisely correct. It also means that it is not necessary to worry too much if the first two or three letters in the type number of a supplied device are not what you were expecting. If you require an MC1458CP but are supplied with a CA1458E there is no need to panic.

End game

There is an obvious problem with this example, where the third part of the type number is also different, and it is only the number in the middle that is common to both. The final part of the type number normally consists of one or two letters, and indicates the type of package.

Most designs for the home constructor are based on integrated circuits (ICs) that are contained in DIL (dual in-line) plastic encapsulations. Dual in-line simply means that the component has two rows of pins. Although there is a degree of standardisation here, not all manufacturers use the same suffix letters for a given package type.

In the type number example provided previously, the suffix letters were 'CP' and 'E'. In the first type number, the C and P respectively indicate a dual in-line package and that it is made from plastic. The single letter E in the second type number means exactly the same thing. Other manufacturers use different suffixes for this type of encapsulation, such as 'C', 'CN', 'N', 'CS', 'P', and 'G'. No doubt there are many other alternatives.

Until relatively recently there were few devices that were offered to amateur consumers in more than one case style. These days, you may occasionally buy components from suppliers that are primarily selling components to the electronics industry rather than the hobbyist. Consequently, some chips will be listed with two or more case styles. Even some suppliers to the amateur market include a few surface-mount versions. It is, therefore, necessary to pay a little more attention to the suffix than it was in the past. The component catalogue should make it perfectly clear which particular type of case each version has.

Manufacturers try to avoid duplication of the middle part of a type number, which is

the 'real' type number, so that the confusion this could cause is avoided. The convention is for the basic type number to be from three to five characters long, and, with a few exceptions, it consists entirely of numerals. While it is not inconceivable that you could find a semiconductor that is totally different to the chip you require, but has the same basic type number, the chances of this happening in practice are remote.




Fig.1. This voltage regulator is a 78L15, which means that it provides an output potential of +15V at output currents of up to 0.1 amps (100mA)

Even so, it is prudent to look at the descriptions of semiconductors to check that they match up with the required device. This might just prevent an expensive error. If a design requires an audio amplifier but the device you find with the right basic type number is a logic chip, clearly you have not found the right device. Some further searching through the catalogue is required.

Up to speed

Some computer chips have a two section suffix that indicates the speed of the device as well as its case type. This is most likely to be encountered with memory and microprocessor chips, including microcontrollers. As a practical example, the PIC16C77-04/P microcontroller is also available as the PIC16C77-20/P. The 'P' part of the suffix in both cases indicates that the component is contained in a standard plastic DIL encapsulation. The '04' and the '20' parts of the suffix indicate that the chips will operate at maximum clock frequencies of 4MHz and 20MHz respectively.

In general, a chip having a higher speed rating will work perfectly well in place of a slower version, but this is almost certain to mean using a significantly more expensive chip where a cheaper type would suffice. Obviously, a slower chip is unlikely to work in a circuit that requires a faster version. Even where this type of substitution seems to be successful it is likely that there will be reliability issues.

Voltage regulators

Component lists and catalogues often refer to many of the more common voltage regulators under their basic type numbers, with no prefixes or suffixes being used. This is presumably due to these devices being manufactured by numerous companies, which results in real-world devices having a bewildering assortment of full type numbers. Voltage regulator type numbers are sometimes abandoned altogether,

A component having 78L15 as its type number (Fig.1.) would therefore be a 15V, 100mA *positive* voltage regulator, and one having 79M05 as the type number would be a five volt 500mA *negative* regulator.

Logical numbering

The situation is similar with the 74 series TTL logic integrated circuits. These devices are mostly sold under a basic type number, which starts with '74' and then has a two or three digit serial number.

The original range of devices is now largely obsolete, but there have been numerous ranges of improved devices over the years. Most of these have now been superseded as well, or simply failed to gain acceptance in the first place.

What is essentially the same numbering system has been retained for the improved devices, but some letters are added between the '74' and the serial number to denote which family the device comes from. This is 'LS' for low-power Schottky, 'HC' for high-speed CMOS and 'HCT' for the high-speed CMOS devices that operate at normal TTL voltage levels. The original 7410 is therefore available as the 74LS10, the 74HC10, and the 74HCT10.

When dealing with TTL integrated circuits it is as well to bear in mind that compatibility between the various families is not very good. They do not even share a common supply voltage range. Using a device from the wrong family is unlikely to provide satisfactory results, and in some cases it could even result in damage to the substitute device or other components in the circuit. Logic integrated circuits are widely available in dual in-line and surface-mount varieties, so it is important to make sure that the right type is ordered.

The other common type of logic integrated circuit is the CMOS 4000 series. These have a basic type number that is just a serial number, starting at 4000. At one time it was possible to buy the original 'A' suffix CMOS devices and the newer 'B' series chips. The A series have been obsolete for many years now, and all the devices listed in the catalogues are B series chips and have a B suffix to the basic type number. There should be no problem in using B series CMOS components where an old design specifies the use of A series chips.

Pro Electron

Many integrated circuits of European origin use the Pro Electron method of type numbering. This follows what is essentially the standard three section system of type numbering, but in a somewhat modified form. The first three letters do not indicate the manufacturer, but instead give some basic information about the device. The first letter is 'S' for a digital device, 'T' for an analogue device, or 'U' for one that is a combination of the two.

The second letter is a serial letter of no real significance, and the third letter indi-

cates the operating temperature range. The third letter is often 'A', which indicates that the device does not conform to one of the standard temperature ranges. The middle section is the usual type number and the final part is a single letter to indicate the type of package. This is often a 'P', indicating a plastic dual in-line package.

The Pro Electron system is also used for many diodes, transistors, and other small semiconductors. The Pro Electron type numbers for these devices consist of two letters followed by a serial number. The first letter is 'A' for a germanium device, 'B' for a silicon type, or 'C' for one that is constructed from gallium arsenide. The second letter indicates the type of component. As a couple of examples, this letter is 'A' for a small diode and 'C' for a low-power audio transistor. The BC549 is therefore a low-power silicon transistor for audio use.

Some Pro Electron codes include a single-letter suffix, and this is most likely to be encountered with transistors that are graded into gain groups. This normally takes the form of letters from 'A' to 'C' being used to indicate low, medium and high gain groups. It is advisable to use a device in the specified gain group where appropriate. Any device of the right basic type can be used where no gain group is specified, including components that have not been graded and lack the suffix letter.

A booklet in PDF format that fully explains the Pro Electron system, is available at this web address:

http://www.eeca.org/pdf/new_pro_electron_fourteenth_2007_06.pdf

JEDEC

The JEDEC system is the American equivalent of the European Pro Electron type. It is less informative, with only the first digit of the type number having any real significance. This is one less than the number of leads that the device has. The second digit is always 'N', and the rest of the type number is a serial number.

Diodes have two leads, and therefore use 1N**** type numbers, whereas normal three lead transistors have 2N**** type numbers.

Optional extras

Beginners are often confused by the extra markings found on many real-world components. Semiconductors are probably worse than most other types of component in this regard. Any additional markings are unlikely to be of any special significance.

One of the extra markings will usually be the manufacturer's logo, and the country of manufacture might also be included. Any additional numbers can be confusing, but are just things like batch numbers, or the date of manufacture in some strangely coded form. This is typically something like the number of days since the factory was opened.

These irrelevant markings can be confusing at first, but you soon get used to picking out the type number from the extraneous characters.



Fig.2. Most semiconductors have extraneous markings. These two chips are a CMOS 4001B (top) and a DAC0801LCN digital-to-analogue converter chip

with the maximum voltage and current ratings being specified instead.

As with other devices, voltage regulators are easier to deal with once you understand the way in which the basic type numbering operates. Regulators for use with *positive* supplies have a type number that starts '78', while those for operation with *negative* supplies have type numbers that commence with '79'. With a device that can operate at up to one amp, the rest of the type number is two digits that indicate the output voltage.

For instance, the two digits are '05' for a 5V regulator and '12' for a 12V type. A 12V positive regulator is therefore a 7812, and a 5V negative type is a 7905. The standard voltage regulators are available with about half a dozen or so output voltages from five to 30V.

Regulators having maximum operating currents of other than one amp are available, and with these a letter inserted in the middle of the type number indicates the current rating. There are only three commonly available alternatives to the standard devices, as follows:

Letter	Current Rating
L	0.1 amps (100mA)
M	0.5 amps (500mA)
S	2 amps



EPE EVERYDAY PRACTICAL ELECTRONICS



UK readers you can
SAVE 54p
on every issue of *EPE*

How would you like to pay £2.96 instead of £3.50 for your copy of *EPE*?
Well you can – just take out a one year subscription and save 54p an issue,
or £6.50 over the year

You can even **save 75p an issue** if you subscribe for two years
– a total saving of **£18.00**

Overseas rates also represent exceptional value

You also:

- Avoid any cover price increase for the duration of your subscription
- Get your magazine delivered to your door each month
- Ensure your copy, even if the newsagents sell out

Order by phone or fax with a credit card or by post with a cheque or postal order, or
buy on-line from **www.epemag.co.uk** (click on “Subscribe Now”)

EPE SUBSCRIPTION PRICES

Subscriptions for delivery direct to any address in the UK:
6 months £18.75, 12 months £35.50, two years £66; Overseas:
6 months £21.75 standard air service or £30.75 express airmail,
12 months £41.50 standard air service or £59.50 express airmail,
24 months £78 standard air service or £114 express airmail.
Cheques or bank drafts (in **£ sterling only**) payable to *Everyday Practical Electronics* and sent to *EPE* Subs. Dept., Wimborne Publishing Ltd., Sequoia House, 398a Ringwood Road, Ferndown, Dorset BH22 9AU. Tel: 01202 873872. Fax: 01202 874562. Email: subs@epemag.wimborne.co.uk. Also via the Web at <http://www.epemag.co.uk>. Subscriptions start with the next available issue. We accept MasterCard, Amex, Diners Club, Maestro or Visa. (For past issues see the *Back Issues* page.)

ONLINE SUBSCRIPTIONS

Online subscriptions, for downloading the magazine via the Internet, \$15.99US (approx. £9.00) for one year available from **www.epemag.com**.

USA/CANADA SUBSCRIPTIONS

To subscribe to *EPE* from the USA or Canada please telephone Express Mag toll free on 1877 363-1310 and have your credit card details ready. Or fax (514) 355 3332 or write to Express Mag, PO Box 2769, Plattsburgh, NY 12901-0239 or Express Mag, 8155 Larrey Street, Anjou, Quebec, H1J 2L5.
Email address: expsmag@expressmag.com.

Web site: www.expressmag.com.

USA price \$60(US) per annum, Canada price \$97(Can) per annum – 12 issues per year.

Everyday Practical Electronics, periodicals pending, ISSN 0262 3617 is published twelve times a year by Wimborne Publishing Ltd., USA agent USACAN at 1320 Route 9, Champlain, NY 12919. Subscription price in US \$60(US) per annum. Periodicals postage paid at Champlain NY and at additional mailing offices. POSTMASTER: Send USA and Canada address changes to *Everyday Practical Electronics*, c/o Express Mag., PO Box 2769, Plattsburgh, NY, USA 12901-0239.

SUBSCRIPTION ORDER FORM



- ☐ 6 Months: UK £18.75, Overseas £21.75 (standard air service), £30.75 (express airmail)
☐ 1 Year: UK £35.50, Overseas £41.50 (standard air service) £59.50 (express airmail)
☐ 2 Years: UK £66.00, Overseas £78.00 (standard air service) £114 (express airmail)
To: *Everyday Practical Electronics*,
Wimborne Publishing Ltd., Sequoia House,
398a Ringwood Road, Ferndown, Dorset BH22 9AU
Tel: 01202 873872 Fax: 01202 874562
E-mail: subs@epemag.wimborne.co.uk

I enclose payment of £ (cheque/PO in £ sterling only), payable to *Everyday Practical Electronics*

My card number is:
Please print clearly, and check that you have the number correct

Signature

Card Security Code (The last 3 digits on or just under the signature strip)

Card Ex. Date Maestro Issue No.

Name

Address

Post code Tel.

Subscriptions can only start with the next available issue.

Check your DMM's accuracy with this:

MiniCal 5V Meter Calibration Standard

How accurate is your digital multimeter? Find out with this simple yet accurate DC voltage reference. If your meter fails the grade, the reference can be used as the calibration standard too. And as a bonus, we've thrown in a crystal-locked frequency reference which doubles as a crystal checker.



By BARRY HUBBLE

RECENTLY, THE NEED arose to recalibrate an expensive digital multimeter. As the job seemed quite straightforward, I decided to tackle it myself. Like most hobbyists, I don't have access to the high-accuracy voltage standards used in calibration labs. Nevertheless, I came up with a scheme that I thought would be accurate enough for general hobbyist work.

By hooking up five multimeters and two panel meters to a voltage divider across a battery, I figured that the mean reading should serve as a reasonable 'standard'. However, I was amazed to see that no two meters read the same and the range of values was much greater than I had anticipated.

Although the readings were probably within the specs for each meter, it was a sobering demonstration. In the

absence of anything better, I calibrated my upmarket digital meter to the mean value but was determined to find a more accurate method that would give me some confidence.

MiniCal solution

The Maxim range of IC voltage references proved ideal for this purpose. In particular, the MAX6350 +5V DC reference boasts a very impressive untrimmed accuracy of $\pm 0.02\%$, with an extremely low temperature coefficient of $0.5\text{ppm}/^\circ\text{C}$. Generally, voltmeters are calibrated on their lowest DC range (200mV for 3.5-digit meters). The 'MiniCal', as this new project is called, divides down the MAX6350's +5V output to generate a 192.3mV reference.

In addition, the board includes a crystal-locked oscillator for checking

meters, oscilloscopes and the like. The frequency of the oscillator is determined by crystal selection.

How it works

Fig.1 shows that the circuit consists of two completely separate sections.

With slide switch S1 in the lefthand position, battery power is applied to the oscillator section. This is comprised of a basic Colpitts oscillator.

The circuit can also be used as a crystal checker if so desired.

ABOVE: our Tektronix 4.5-digit meter is pretty much spot on, especially when the 0.02% accuracy of the MiniCal voltage reference is considered. Other (cheaper) meters might not be as accurate.

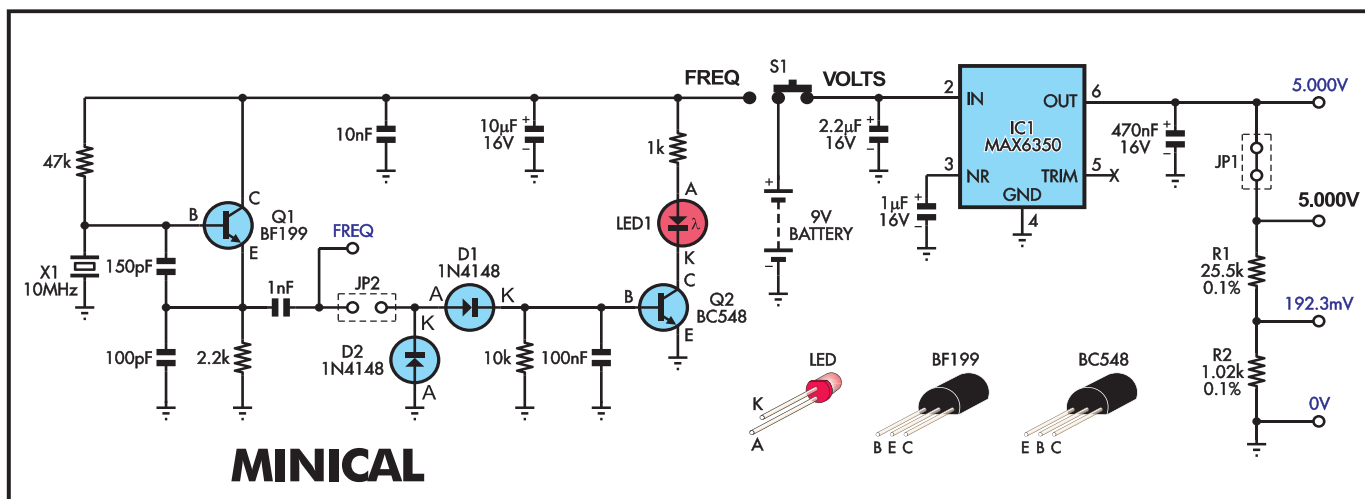


Fig.1: the MiniCal consists of independent oscillator and voltage reference circuits. To minimise noise on the voltage reference, only one of the circuits can be powered at a time, selectable via slide switch S1.

Reproduced by arrangement with
SILICON CHIP magazine 2007.
www.siliconchip.com.au

Crystal X1, the 150pF capacitor between Q1's base and emitter, and the 100pF capacitor to ground (0V) together form the feedback network. The output from Q1's emitter is AC-coupled via a 1nF capacitor to the 'FREQ' test pin.

Although we've specified a 10MHz crystal for X1, the circuit should work with values from 1MHz to at least 21MHz without modification.

The remaining circuitry connected to Q1's emitter performs the crystal 'go/no go' function. Diodes D1, D2 and the 100nF capacitor rectify and filter the AC signal from the emitter. The resultant DC voltage is applied to the base of Q2, switching it on and lighting the 'OK' LED whenever oscillation is present.

Voltage reference

With switch S1 in the righthand position, the voltage reference section of the circuit is powered. This section is very simple and consists of only a voltage reference IC, three capacitors and two resistors.

The MAX6350 (IC1) can operate with an input range of 8-36V, providing an untrimmed output of 5V $\pm 0.02\%$ (4.999V – 5.001V). Small tantalum capacitors on the input, output and 'NR' (Noise Reduction) pins reduce circuit noise to just 3.0 μ Vp/p (typical) in the 0.1Hz to 10Hz spectrum. Battery-powered operation ensures that this is not degraded by external (conducted) noise sources.

Note: the MAX6350 is available in

both 8-pin DIP and SO (surface mount) packages. The PC board design accommodates both package styles. We expect that most constructors will opt for the surface mount device, as it is cheaper and easier to obtain.

Resistors R1 and R2 divide down the MAX6350's +5V output to obtain the 192.3mV calibration voltage. At a minimum, these resistors need to be 0.1% types (see parts list) to achieve the specified 0.2% voltage tolerance.

As you can see, the use of 0.1% resistors degrades circuit performance somewhat. However, the result is a good compromise between accuracy and cost, and is sufficient for meter checking. If you want to use the MiniCal for calibration, then you will need to upgrade to tighter-tolerance resistors in order to meet the basic accuracy specs of your instrument.

Two alternatives for R1 and R2 are shown in the parts list. The 0.01% resistor pair gives a $\pm 0.04\%$ tolerance

on the 192.3mV output, but they are expensive. Alternatively, you can install the 0.05% 25:1 divider network for a tolerance of about 0.1%, and at a much lower cost.

Note: the 25:1 divider network consists of two 0.1% resistors (1k Ω and 25k Ω) with a ratio accuracy of 0.05%. The device is supplied in a 3-pin surface-mount (SOT-23) package.

So why did we choose an odd calibration voltage of 192.3mV instead of a nice round figure? Well, it was simply a convenient choice using available resistor values. Other division ratios could be used but for best results the reference voltage must be close to (but not exceeding) 200mV.

Construction

All parts mount on a single PC board coded 622 – see Fig.2. If you have surface-mount devices for IC1 and/or R1 and R2, these should be installed first (see Fig.3). You'll need a temperature-controlled soldering iron with a fine chisel tip and small-gauge solder for the job. A bright light, magnifying glass and 0.76mm desoldering braid ('Solder-Wick' size #00) will also prove useful.

Next, on the top side of the board (see Fig.2), install all components in order of height, starting with the wire link, resistors and diodes (D1 and D2). Obviously, if you've mounted the R1/R2 divider on the bottom side, then you shouldn't install anything in the R1 and R2 positions on this side!

Note that all the tantalum capacitors are polarised devices and must be inserted with their positive leads aligned with

Main Features

- 5.000V $\pm 0.02\%$ voltage standard
- 192.3mV $\pm 0.2\%$ voltage standard (optional $\pm 0.1\%$ or $\pm 0.04\%$)
- Two $\pm 0.1\%$ resistor standards (optional $\pm 0.01\%$)
- Crystal-locked frequency reference
- Crystal checker

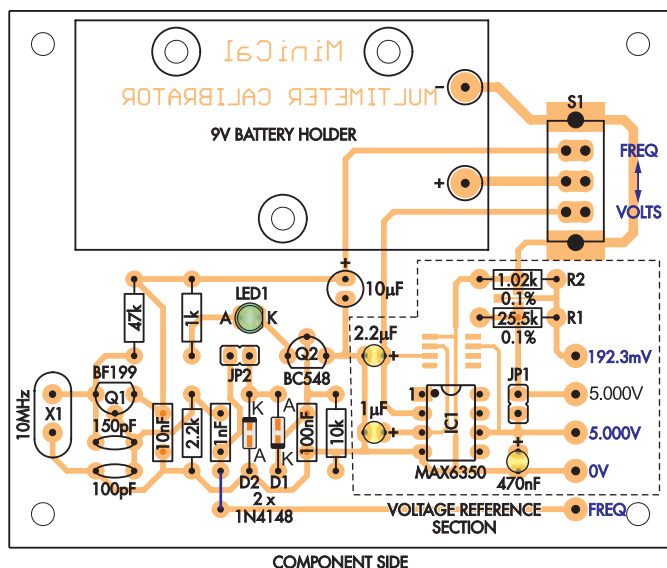


Fig.2: follow this diagram closely when assembling the board. Take care with the orientation of the diodes (D1 and D2) and tantalum capacitors. Note: this final version of the PC board differs slightly from the early version shown in the photographs

the '+' symbol marked on the overlay.

Install the battery holder last of all. It should be fixed to the PC board with No.4 x 6mm self-tapping screws before soldering.

To complete the job, attach small stick-on rubber feet to the underside of the PC board to protect the assembly as well as your desktop.

Operation

Due to the expected intermittent use of the MiniCal, a power switch has not been included. Simply plug

in a battery and use the slide switch to select between the oscillator function (FREQ) or voltage reference function (VOLTS). Note that the battery voltage must be at least 8V for correct operation of the reference IC.

When measuring the oscillator frequency, the crystal checker function must be disabled by removing the jumper from JP2. This is necessary because the checker circuit loads the oscillator, reducing the signal on the FREQ test pin below the sensitivity level of most multimeters.

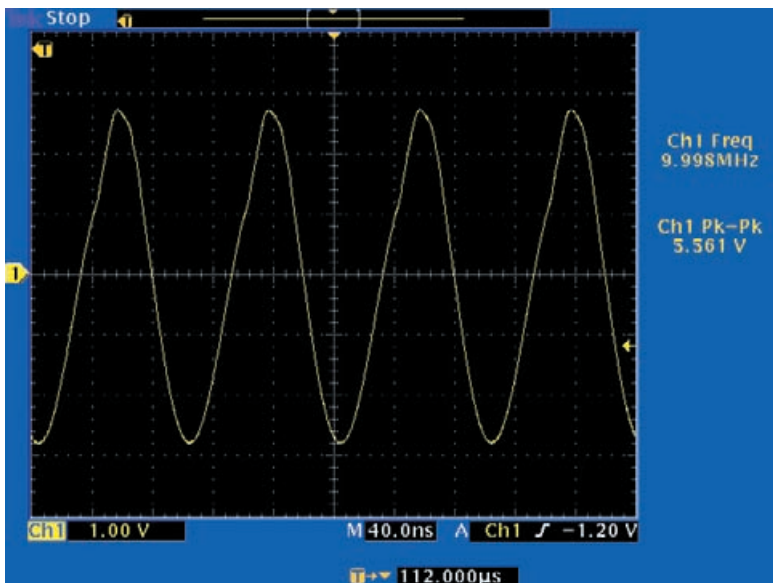


Fig.4: this oscilloscope shot shows the signal on the 'FREQ' test pin with a 10MHz crystal installed. Fig.5 (right) shows the full-size etching pattern for the PC board.

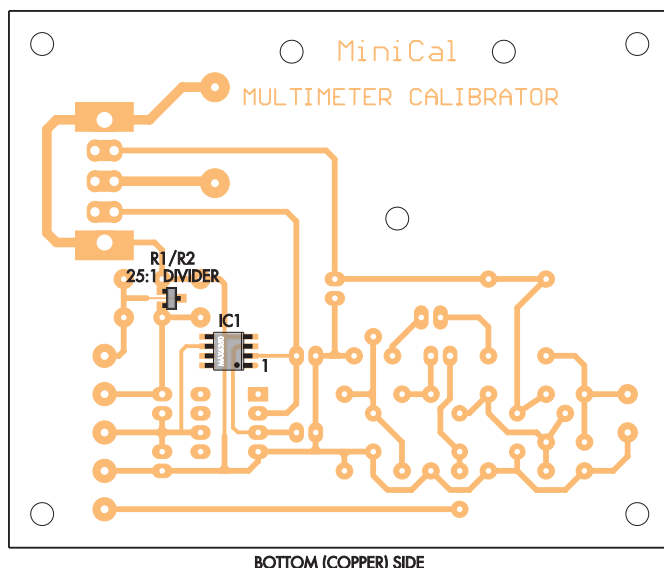



Fig.3: the PC board design can accommodate both conventional (DIP-8) and surface-mount (SO-8) package types for IC1. If you have the SO-8 type, then mount it on the copper side of the board as shown here. The optional 25:1 resistor network (R1/R2) also goes on this side.

Follow the instructions provided with your multimeter regarding calibration. In general, most multimeters should be calibrated on their lowest (basic) range, which is normally 200mV for 3.5 digit models.

As described earlier, accuracy will be about $\pm 0.2\%$ using $\pm 0.1\%$ resistors for R1 and R2. This figure is good enough for many general-purpose instruments, which typically specify an accuracy of $\pm 0.25\%$ at best. Note that calibration instructions usually specify a standard of $\pm 0.1\%$ or better.



Parts List – MiniCal

- 1 PC board, code 622, available from the *EPE PCB Service*, size 71mm x 88mm
- 1 10MHz crystal (X1) (user select, see text)
- 1 3mm green LED (LED1)
- 5 PC board pins (stakes)
- 2 2-way 2.54mm SIL headers (JP1, JP2)
- 2 jumper shunts
- 1 miniature DPDT PC-mount slide switch
- 1 9V PC-mount battery holder
- 3 No.4 x 6mm self-tapping screws
- 4 small stick-on rubber feet
- 1 9V battery

Semiconductors

- 1 MAX6350CPA (DIP) or MAX-6350CSA (SMD) voltage reference (IC1) (Farnell)
- 1 BF199 NPN RF transistor (Q1)
- 1 BC548 NPN transistor (Q2)
- 2 1N4148 diodes (D1, D2)

Capacitors


- 1 10 μ F 16V tantalum
- 1 2.2 μ F 16V tantalum
- 1 1 μ F 16V tantalum
- 1 470nF 16V tantalum
- 1 100nF 63V MKT polyester
- 1 10nF 63V MKT polyester
- 1 1nF 63V MKT polyester
- 1 150pF ceramic disc
- 1 100pF ceramic disc

Resistors (0.25W, 1%)

- | | |
|--|-----------------|
| 1 47k Ω | 1 2.2k Ω |
| 1 10k Ω | 1 1k Ω |
| 1 25.5k Ω 0.1% (R1) (Farnell) | |
| 1 1.02k Ω 0.1% (R2) (Farnell) | |
| -OR- | |
| 1 25:1 0.05% resistor network, Vishay MPM series (Farnell) | |
| -OR- | |
| 1 25k Ω 0.01%, Vishay S102J series (Farnell) | |
| 1 1k Ω 0.01%, Vishay S102J series (Farnell) | |

Table 1: Capacitor Codes

Value	μ F Code	EIA Code	IEC Code
470nF	0.47 μ F	474	470n
100nF	0.1 μ F	104	100n
10nF	.01 μ F	103	10n
1nF	.001 μ F	102	1n
150pF	—	151	150p
100pF	—	101	100p



Calibration is normally only applicable to the basic range, with all other ranges depending on that calibration. The 5V output and 0.1% resistors should therefore only be used to check the accuracy of your meter, not to calibrate it. Note that, in use, the jumper shunt (on JP1) must be removed before measuring the 0.1% resistor values.

Note also that some meters may require special tools and/or knowledge for successful calibration. When in doubt, read the (service) manual first!

Meter loading effects

A resistive divider was chosen to generate the millivolt source because it's simple and requires no adjustment. However, the down side to this simplicity is that the meter's input impedance loads the divider network and therefore reduces the reference accuracy.

For example, when a meter with a 10M Ω input impedance is connected, the reference voltage will fall by about 0.02mV. This corresponds to a 0.01% reduction in accuracy. Assuming you know your meter's input impedance, the loading effect can easily be factored into the calibration where maximum accuracy is required.

Further reading

Detailed technical information on the MAX6350 voltage reference IC can be downloaded from the Maxim web site at www.maxim-ic.com *EPE*

ANDRE LAMOTHE'S

XGAMESTATION

LEARN STEP-BY-STEP HOW TO DESIGN AND BUILD YOUR OWN VIDEO GAME CONSOLE!

Design inspired by the Atari 800/2600, Sinclair ZX Spectrum, Apple II & Commodore 64!

Complete Package

eBook

Integrated IDE

SX52 CPU 80 MIPS!

OPEN SOURCE!

FEATURES:

- Great for Hobbyists AND Students!
- Complete Software Development Kit!
- eBook on Designing the XGS Console!
- Parallax SX-Key Compatible!
- Fully Assembled XGS Micro Edition Unit!
- The Fun Way to Learn Embedded Systems!

PAL & NTSC COMPATIBLE!

WWW.XGAMESTATION.COM

SUPPORT@NURVE.NET | PH 925.736.209B(USA)

Ingenuity Unlimited

WIN A PICO PC-BASED
OSCILLOSCOPE WORTH £799

- 200MHz Analogue Bandwidth Dual Channel Storage Oscilloscope
- Spectrum Analyser
- Multimeter
- Frequency Meter
- USB Interface.

If you have a novel circuit idea which would be of use to other readers then a Pico Technology PC-based oscilloscope could be yours.

Every 12 months, Pico Technology will be awarding a PicoScope 3206 digital storage oscilloscope for the best IU submission. In addition a PicoScope 2105 Handheld 'Scope worth £199 will be presented to the runner up.

Our regular round-up of readers' own circuits. We pay between £10 and £50 for all material published, depending on length and technical merit. We're looking for novel applications and circuit designs, not simply mechanical, electrical or software ideas. Ideas *must be the reader's own work* and **must not have been published or submitted for publication elsewhere**. The circuits shown have NOT been proven by us. *Ingenuity Unlimited* is open to ALL abilities, but items for consideration in this column should be typed or word-processed, with a brief circuit description (between 100 and 500 words maximum) and include a full circuit diagram showing all component values. **Please draw all circuit schematics as clearly as possible**. Send your circuit ideas to: *Ingenuity Unlimited*, Wimborne Publishing Ltd., Sequoia House, 398a Ringwood Road, Ferndown, Dorset BH22 9AU. (We **do not** accept submissions for IU via email.) Your ideas could earn you some cash and a prize!



Micropower Cell Monitor – *Striving For Equality*

THE circuit shown in Fig.1 will monitor the charge on all the cells in a battery pack. A series-connected battery pack of several NiCad or NiMh cells is very commonly used to provide power for portable projects and devices. However, all cells are not created equal and this can lead eventually to the premature failure of one or more of the cells in the pack.

Consider the situation when one of the cells in a pack has been depleted before the others (even a tiny capacity variation between cells will inevitably lead to this situation). Although the overall pack voltage may still be sufficient for the powered device, the single depleted cell can actually become reverse biased (reversed charged) while load current is still being drawn.

The reverse bias on the cell will actually lead to considerable cell damage, which then leads to even earlier depletion after the next charge cycle. So even a tiny capacity shortage in one of the cells can lead to a massive capacity shortage after a few cycles.

The circuit here monitors the voltage across every cell in a pack (you simply need

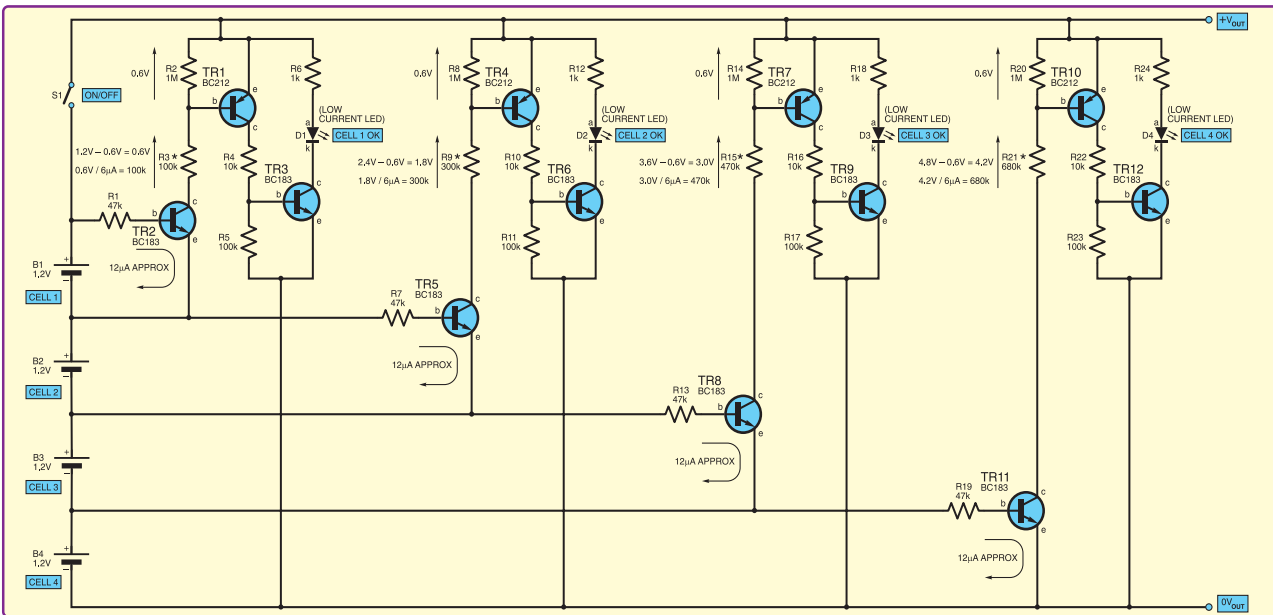


Fig.1. Circuit diagram for the Micropower Cell Monitor

access to all the cell connections). If any cell's voltage drops below around 0.7V, then its associated transistor will turn off and the cell indicator LED will extinguish.

So, while you're using your equipment you can ensure that you switch off and recharge as soon as any of the status LEDs go out, thereby lengthening the life of the cells. You could even choose to replace just the cell(s) that are failing too soon, thereby increasing the capacity of the whole pack without replacing the whole lot in one go.

While the main power switch is off, the current consumption of the cell monitor is extremely low (around 20µA per cell). Switched on, there is, of course, the extra consumption of the individual LEDs. The cell consumption (of the circuit) is in fact less than 1/20th of a standard AA cell's self-discharge. If you ensure that the circuit is active only when the equipment itself is powered, then it is highly likely

that the monitor's consumption will be only a tiny proportion of the equipment's consumption.

The circuit is easily adapted for any number of cells, just add extra stages. To keep current consumption at a minimum though, please ensure that the resistors marked with an asterisk are calculated to allow approximately 6µA of current through each at the voltage that appears across it.

For example, resistors R3, R9, R15 and R21 drop a progressively higher voltage as more of the battery pack's cells are seen by each successive stage. Each stage sees approximately 1.2V more than the preceding stage. So, roughly, add 200kΩ to the resistor value for each extra stage. It doesn't need to be exact at all, it's just a current consumption issue, little to do with thresholds.

Finally, remember that each stage needs connections to the supply rails (+V_{out} and

INGENUITY UNLIMITED

BE INTERACTIVE

IU is your forum where you can offer other readers the benefit of your Ingenuity. Share those ideas, earn some cash and possibly a prize.



GND (0V) rails) as well as the connections to each cell.

Jez Siddons
Chapel-en-le-Frith

10V to 15V Battery Voltmeter – *Currently Charging*

AFTER suffering alternator problems on my car, I decided to build a limited range meter to monitor its battery voltage from the cigar lighter socket. The circuit is shown in Fig.2.

The circuit uses a 741 op amp, IC1, as a voltage comparator, with its non-inverting input (pin 3) held at a constant 6.8V by Zener diode D1 and resistor R1. Its inverting input (pin 2) is fed by a fraction of the battery voltage via potential divider network R2/VR1/R3. This divider network holds the inverting input at 68% of the input voltage, so when the battery voltage is 10V, both inputs are at the same voltage and no current flows through the meter. VR1 should be




Fig.2. Circuit diagram for the 10V to 15V Battery Voltmeter

adjusted to ensure that the meter reads zero at 10V input voltage.

As the battery voltage increases from 10V, the op amp input voltages start to differ, so an output current flows through the meter, and back through VR2 into the potential divider network to hold the inverting input at 6.8V. The value of this current can be adjusted by VR2 so that the meter full scale deflection can be set to 15V.

The 10nF capacitor, C1, is needed to ensure the stability of the op amp, and omitting it could result in false readings.

It was found that with a good car alternator, the battery voltage should be just over 14V, and should fall only very slightly when the headlights are switched on.

P.A. Tomlinson,
East Yorkshire

EPE BINDERS

KEEP YOUR MAGAZINES SAFE – RING US NOW!

This ring binder uses a special system to allow the issues to be easily removed and re-inserted without any damage. A nylon strip slips over each issue and this passes over the four rings in the binder, thus holding the magazine in place.

The binders are finished in hard-wearing royal blue p.v.c. with the magazine logo in gold on the spine. They will keep your issues neat and tidy but allow you to remove them for use easily.

The price is £7.95 plus £3.50 post and packing. If you order more than one binder add £1 postage for each binder after the initial £3.50 postage charge (overseas readers the postage is £6.00 each to everywhere except Australia and Papua New Guinea which costs £10.50 each).

Send your payment in £'s sterling cheque or PO (Overseas readers send £ sterling bank draft, or cheque drawn on a UK bank or pay by card), to **Everyday Practical Electronics, Wimborne Publishing Ltd, Sequoia House, 398a Ringwood Road, Ferndown, Dorset BH22 9AU.**

Tel: 01202 873872. Fax: 01202 874562.

E-mail: editorial@epemag.wimborne.co.uk.

Web site: <http://www.epemag.co.uk>

Order on-line from:

www.epemag.wimborne.co.uk/shopdoor.htm

We also accept card payments. Mastercard, Visa, Amex, Diners Club or Switch (minimum card order £5). Send your card number and card expiry date plus Switch Issue No. with your order.



Professional possibilities with your digital camera!



By Terry de Vaux-Balbirnie BSc

Digi-Flash Slave

DIGITAL cameras have advanced enormously over the past few years. The author bought his first one in 1998 – a one megapixel unit costing almost £600. Now, digital cameras having a resolution of six megapixels or more are available at a fraction of that price.

Although 'point and shoot' digital cameras may be fine for outdoor shots, users often report that photographs taken indoors are disappointing. The tiny built-in flash units are barely adequate for anything but snapping at parties or as 'fill-in' illumination for outdoor work. For serious indoor photography, more light is needed and preferably provided by more than one flash unit.

Convenience

Even so, the flash built into the camera is convenient and, for many users, this is the overriding consideration. Any drawbacks will probably have been dismissed as inevitable. One obvious problem is the shadow cast

on a surface behind the subject. This appears larger than the subject itself and displaced from it slightly.

Because a flash tube is a small source of illumination it gives a dark, sharp-edged shadow compared with that provided by a larger light source. Professional photographers 'soften' shadows by rotating the flash head so that the light reflects from a large surface such as a wall, ceiling or flash umbrella. This, in effect, makes the light source larger. This cannot be done with a fixed flash that always points to the front.

Another problem with the built-in flash is that the subject's pupils have a crimson appearance. This 'red-eye'

is caused by light reflecting from the retina, which is rich in blood vessels. Since the light reflects straight back to the camera's lens, red-eye is a particular problem with point-and-shoot equipment.

It is more evident in dim light because the iris is wide open. On many



cameras, there is a 'red-eye correction lamp' which switches on just before the flash fires. This causes the iris to close a little and so makes the effect less noticeable. This is not usually very effective and red-eye can only be properly eliminated by placing the light source, or sources, some distance to the side of the camera's lens.

Two is better than one

Two or more flash units can give much better results than the built-in one. One arrangement is shown in Fig.1. Flash units A and B illuminate each side of the subject. If they are equally powerful, one should be placed a little closer than the other. This avoids the 'flat' result obtained when the subject is evenly illuminated.

Reflectors are used to soften the shadows and a further flash might be used to highlight the hair, for example. Since the flash units are placed to the side of the camera, red-eye is eliminated. Even a single auxiliary flash can give good results if the light is reflected from a large surface placed behind the camera position.

Suitable flash units for this type of work are not expensive. Medium-powered battery operated equipment is usually adequate and may be bought on an on-line auction for a few pounds. Reflecting surfaces may be home-made using sheets of foam polystyrene.

It will be best if the camera can be switched to 'manual' to allow the exposure to be adjusted for the extra light. It would be a good idea to check that this is possible with your equipment. However, if it cannot be done and the camera is used on 'auto-automatic', you might be able to compensate for exposure errors using imaging software. More will be said about this later.

Master and slave

The circuit described here is for a digital camera flash slave unit. Its purpose is to pick up light from the camera's flash (the master) and use it to trigger one or more remote units (slaves). This overcomes the problem of there being no socket on the camera to connect auxiliary flash units. Also, and very importantly, there are no connections made to the camera so there is no possibility of a high voltage pulse from a slave flash damaging it.




Fig.1: Two or more flash slaves can produce better results than a single built-in unit

It is important to note that a flash slave unit made for a traditional outfit is unlikely to work with a digital camera. This is because many popular models give two flashes when taking a photograph rather than one. The 'pre-flash' can have several functions, such as to set the intensity of the main flash and to adjust the white balance. Since a conventional circuit would trigger the slave units on the pre-flash, they would fire too early and fail to synchronise with the shutter.

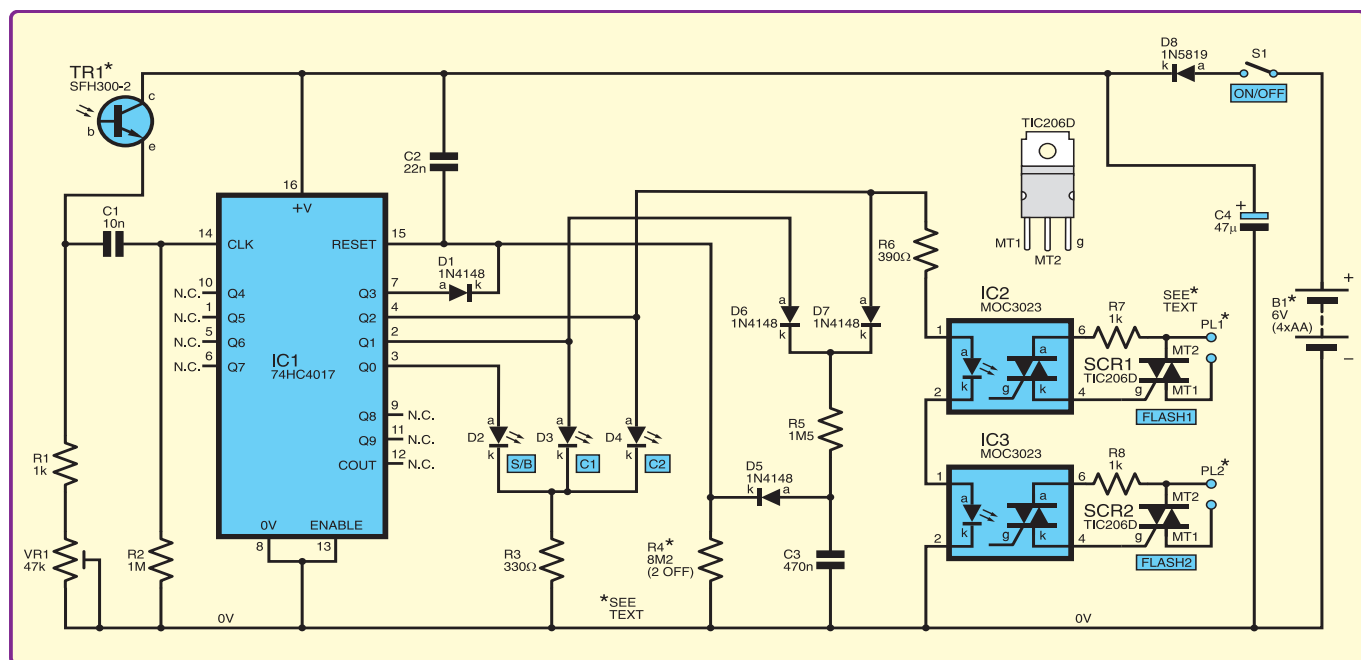


Fig.2: Complete circuit diagram for the Digi-Flash Slave. Optical isolation provided by IC2 and IC3 prevents any high voltage pulse destroying semiconductors in the circuit.

Initial checks

Some digital cameras do not use a preflash so you should find out before proceeding. However, if your camera is a Canon, Olympus, Epson or Nikon it is most likely to use one. You can usually see the double flash if you observe from a distance. If not, check with the manufacturer. Find out also that your camera gives only *one* preflash – there are some professional Nikon and Olympus single-lens reflex cameras (and probably others) that use two.

Slave flash circuits designed for a digital camera must therefore ignore any preflash and there are two possible approaches. One is to count the flashes and disregard the first one. The other is to use a time delay so that, when the preflash is detected, the slave units trigger a little later so that they synchronise. This circuit is of the first type.

Circuit description

The complete circuit for the Digi-Flash Slave is shown in Fig.2. Current is drawn from the 6V battery pack, B1, via on-off switch S1 and diode D8. Capacitor C4 provides a reserve of energy and allows the circuit to operate effectively when the battery is nearing the end of its service life. The diode protects the circuit against connecting the supply the wrong way round. The specified diode is a Schottky device. This minimises the forward voltage drop, which is 'lost' to the circuit.

When phototransistor TR1 detects light from the master flash, the voltage between its collector (c) and emitter (e) falls momentarily. This results in an increased voltage appearing across resistor R1 and preset VR1, and hence at the left-hand side of capacitor C1. The voltage pulse is transferred by C1 to the clock input (pin 14) of decade counter IC1. Resistor R2 maintains this in a normally-low condition which prevents false triggering.

The purpose of C1 is to provide AC coupling between TR1 and IC1, which makes the circuit practically immune from slow changes in light level. The quiescent voltage appearing across R1/VR1 is dependent on the ambient light level and VR1's adjustment – it is set for correct operation at the end of construction.

Counting pulses

Decade counter IC1 has ten outputs (Q0 to Q9) which go high in turn as

Parts List – Digi-Flash Slave

1 PC board, code 625 available from the <i>EPE PCB Service</i> , size 68.5mm x 52mm	1 1N5819 40V 1A Schottky diode (D8)
1 plastic case, size 102mm x 78mm x 38mm	1 SFH300-2 phototransistor or similar – see text (TR1)
4 AA-type alkaline 1.5V cells (B1)	2 TIC206D 400V 4A triacs (SCR1, SCR2)
1 four-cell (AA) battery holder, with leads and clips	1 74HC4017 decode counter – see text (IC1)
1 miniature toggle or rocker switch (S1)	2 MOC3023 opto-coupled triacs (IC2, IC3)
2 3-pin panel type plugs – see text (PL1, PL2)	Capacitors
2 line-sockets to fit 3-pin plugs – see text (SK1, SK2)	1 10nF polyester, 5mm pitch (C1)
2 6-pin DIL sockets (IC2, IC3)	1 22nF polyester, 5mm pitch (C2)
1 16-pin DIL socket (IC1)	1 470nF polyester, 5mm pitch (C3)
1 2-way PCB mounting screw terminal block (TB1)	1 47 μ F radial elect. 16V (C4)
2 plastic PC stand-off insulators or Nylon bolts (2) and nuts (6)	Resistors (0.25W 5% carbon)
Multistrand connecting wire;	1 16M4 (2 x 8M2 – see text) (R4)
small plastic cable tie, see text; solder etc.	1 1M5 (R5)
Semiconductors	1 1M Ω (R2)
4 1N4148 signal diodes (D1, D5 to D7)	3 1k Ω (R1, R7, R8)
2 3mm red LEDs (D3, D4)	1 390 Ω (R6)
1 3mm green LED (D2)	1 330 Ω (R3)
	Potentiometer
	1 47k Ω min. round carbon preset, vertical mounting

each clock pulse arrives. To ensure that output Q0 is high on powering-up, capacitor C2 delivers a momentary pulse to the Reset input (pin 15). The high state of Q0 then operates the green standby (S/B) LED D2. Resistor R3 limits its operating current to some 12mA. As well as signalling that the circuit is ready for use, D2 also behaves as the on-off indicator.

When the preflash and then the main flash arrive, IC1 output Q1 goes high, followed by Q2, which then remains high until the circuit resets – this aspect will be discussed presently. Outputs Q1 and Q2 operate the red LEDs (C1 – Count 1, C2 – Count 2, D3 and D4) respectively.

It will be noted that the group of three LEDs (D2 to D4) share current-limiting resistor, R3. This is acceptable practice because only one of them can be on at the same time. The 'C1' and 'C2' LEDs allow operation of the circuit to be monitored as an aid to setting up. When working correctly, D3 will light briefly on the preflash then D4 will operate.

Trigger happy

Ignore diodes D6, D7 and associated components for the moment. With the arrival of the main flash, the high state of IC1 Q2 directs current through resistor R6 and the LED sections (pins 1 and 2) of optically-coupled triacs, IC2 and IC3, connected in series. Resistor R6 limits the current to a safe working value.

The light from the internal LEDs triggers the associated triacs so that pins 4 and 6 in each device become effectively short-circuited. This, in turn, operates external triacs SCR1 AND SCR2 by directing current into their gate (g), via resistor R7 or R8, as appropriate. Once triggered, the main terminals (MT1 and MT2) of these devices become effectively closed switches.

This allows conduction between the trigger contacts of the slave flash units connected to them through plugs PL1/PL2. Note that the power supply for the triacs is obtained from the 'trigger voltage' that exists across the contacts of the slave flash units. Due

to the bidirectional nature of the triacs, it does not matter which way round a slave flash is connected.

High voltage

The triac section of the circuit will withstand up to 400 volts. This is necessary because the trigger voltage will be in the region of 200V to 300V with older flash guns. Modern equipment uses a much lower trigger voltage but it is thought that most users will use inexpensive older flash units.

Modern flash units using a low trigger voltage may not operate with the arrangement used here. You would need to find out on an individual basis. The optical isolation provided by IC2 and IC3 prevents any high-voltage pulse destroying semiconductors in the circuit.

It will be seen that there are two independent channels to which slave flash units may be connected. It is often possible simply to connect the trigger contacts of several flash units in parallel so, if you wish to use three or more slave flash units, simply try this. However, there might be a clash of polarity which could prevent operation of some units. The separate channels allow 'mixing and matching' of various types of equipment.

Resetting

After a photograph has been taken, it is necessary for IC1 to reset so that the circuit is ready to operate again. This is done by taking a signal from Q1 or Q2 and directing it, via diode D6 or D7 and fixed resistor R5, to capacitor C3. With either of these outputs high, the capacitor will charge and the voltage across it rise. The reason for considering Q1 is because there is a chance that only one flash will be detected (possibly due to poor setting-up). If IC1 did not reset, it would 'freeze' with pin 2 high.

The voltage developed across C3 is applied to IC1's reset input, via diode D5. When it reaches a sufficient value (which takes less than one second), the reset pin will interpret the signal as 'high' and the counter will reset. Output Q0 then reverts to high and the S/B LED will operate again.

Resistor R4 provides a discharge path for capacitor C3 and maintains the reset input in a normally low condition which prevents false resetting. Note that, in normal operation, Q1 will only go high momentarily. This will not reset IC1 because the voltage across C3 will not rise significantly in the available time.

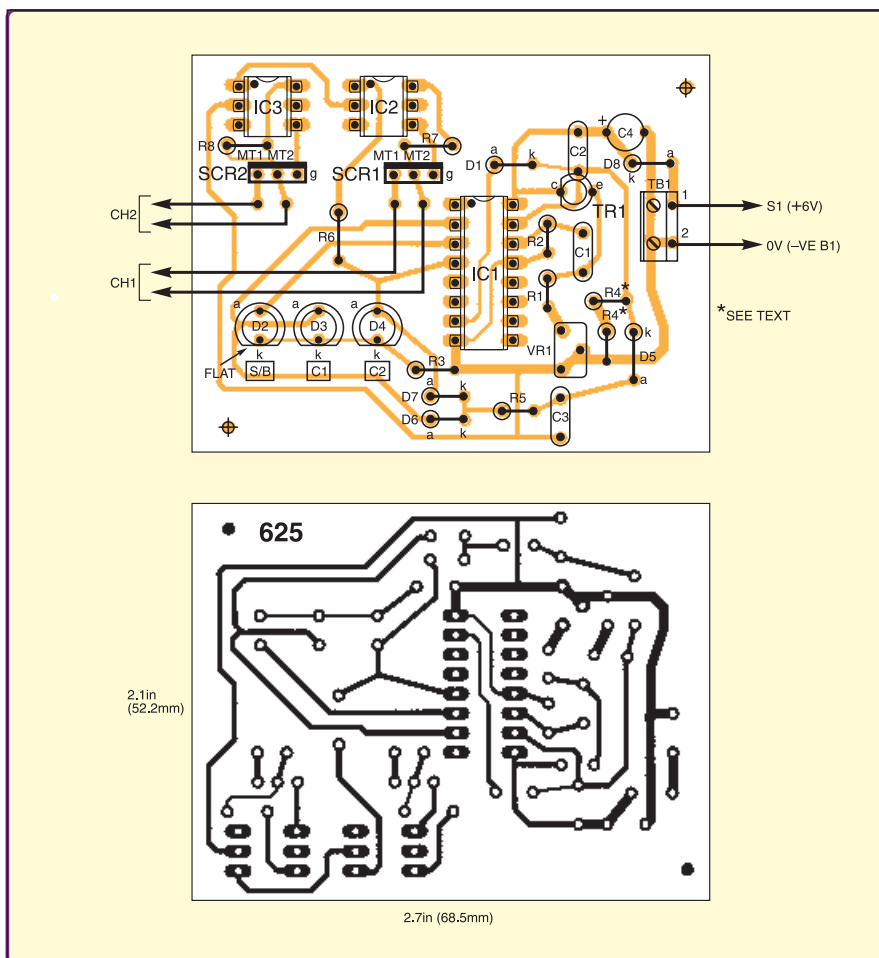


Fig.3: Printed circuit board component layout and full-size copper foil master. Note the phototransistor TR1's package is the same as an LED, i.e. 'flat' equals collector pin

There is a possibility that IC1 could count '3' (perhaps due to an additional flash of light being picked up from another source). If three flashes were detected in rapid succession, the circuit would not reset on the first or second because the voltage across C3 would not rise sufficiently and IC1 would 'freeze' with Q3 high. If this ever happened, a reset signal would be applied to the reset input via diode D1.

If there is an excessive input of light from the master flash and preset VR1 is poorly adjusted, this can cause instability and IC1 may register a few false counts. This can result in the IC 'locking' with one of the unused outputs high and none of the LEDs operating. If this should happen, switch off for a few seconds to allow capacitor C4 (which maintains the state) to discharge. The circuit will reset when switched on again.

Construction


The device specified for IC1 is the 74HC4017 decode counter. This provides

a higher output current than the 4017B, which is *not suitable*. Various phototransistors could be used for TR1 but make sure the one chosen is not an infra-red device housed in an opaque package. It must respond to visible light.

Construction of the Digi-Flash Slave is based on a single-sided printed circuit board. This board is available from the *EPE PCB Service*, code 625. The component layout and actual size copper master pattern is shown in Fig.3.

Begin construction by drilling the two board mounting holes. Solder the IC sockets and screw terminal block, TB1, in place. Follow with all resistors (including preset VR1) and capacitors. Mount capacitor C4 flat on the circuit panel (see photograph) taking care over its polarity.

Resistor R4 should have a value of 16M Ω approximately. This may comprise two 8.2M Ω units connected in series and two pairs of pads have been provided on the PCB for that purpose (both labelled R4).



Layout of components on the prototype circuit board, the phototransistor is just to the left of the terminal block. Note that the radial electrolytic capacitor (C4), top right, is mounted flat on the PCB, its leads being carefully bent at 90° before inserting on the board.

Taking care over their polarity, add diodes D1 and D5 to D8, also triacs SCR1 and SCR2. Mount the LEDs (D2 to D4) using the entire 25mm length of their end leads so that the tops stand higher than everything else. If the leads are shorter than 25mm, you will need to extend them. Mount phototransistor TR1 so that its top stands a little below that of the LEDs. Note that, with the specified unit, the collector (c) (connected to supply positive) has the shorter lead, also there is a 'flat' on the body next to it.

Insert the ICs into their sockets. However, before handling the pins remove any static charge from your body by touching something which is earthed (for example, a metal water tap) to avoid possible damage.

Boxing up

Decide on suitable positions for the various parts inside the box and hold the PCB in place. Mark through the mounting holes, remove the PCB again and drill these holes through. Mount the PCB temporarily on short stand-off insulators so that the tops of the LEDs are a little higher than the lid of the box. Measure their positions and drill holes in the lid for them to show through.

Drill a further hole in the lid, having a diameter of 4mm approximately, directly above phototransistor TR1's position to allow light from the master flash unit to reach it. Drill holes for on-off switch S1 and for the two panel plugs (PL1 and PL2), which will be used to connect the slave flash units.

It is important that the matching line-sockets (which will be fitted to the flash leads) are of a type where it is impossible to touch the pins. This eliminates the possibility of an unpleasant electric shock if the pins were to be touched when the flash unit was charged. Miniature 3-pin mains-type panel-mounting plugs and line sockets were used in the prototype, with only two of the pins used in each case.

An optional hole drilled in the side of the box would allow easy adjustment to VR1 with the lid of the box in place.

Finishing off

Secure the PCB and, referring to Fig.4, complete the wiring. Provide some strain relief at the switch terminals by securing the wires to its body using a small cable tie (see photograph). With the arrangement shown, the battery holder is held between the side of the box and the circuit panel and needs no further support. **Do not** use a 9V (say, PP3) battery – more than a nominal 6V supply would destroy IC1.

The best way to connect the slave units is to use commercial extension leads. The most common connector used on flash units is a small coaxial type and matching sockets seem to be almost impossible to source. Some flash units do not have an integral connecting lead because they are designed to be attached to the camera's 'hot shoe'. It is then necessary to buy a 'hot shoe adaptor'. All these items may be obtained from a good photo shop. Cut the plug off the extension leads and fit the new line-sockets instead.




Fig.4: Interwiring details to off-board components. Don't forget to use the full length (25mm min) of LEDs D2, D3 and D4 pins/leads so that their tops 'sit' comfortably in the case lid holes provided for them.

Testing

Do not connect any slave flash units yet. Adjust VR1 to approximately two-thirds of its total clockwise travel (as viewed from the right-hand side of the circuit board). This should be approximately correct with the specified phototransistor. Insert the batteries into their holder.

Place the lid of the box in position temporarily and point the unit away from bright sources of light. Switch on – the green S/B LED should operate. Test the circuit using the digital camera's flash. The 'C1' LED should flash briefly followed by 'C2', then the circuit should revert to standby.

Adjust preset VR1 as necessary for reliable operation – under average room lighting, the voltage measured across R1/VR1 (between TR1's emitter and supply negative) was 0.5V approximately in the prototype. Best results are obtained when the ambient light is not too bright and the phototransistor can 'see' the camera's flash directly, but at a distance. Allow ten seconds minimum between tests to allow capacitor C3 to discharge sufficiently.

Synchronisation

When satisfied that the unit is triggering correctly, connect a flash unit to one of the outputs. Include this in a trial photograph and check that it shows as a bright patch of light. This proves that the slave unit has synchronised with the shutter (see photograph). Check that the other



Harsh result using inbuilt flash

output works and check with the other flash units.

With both outputs tested, take some trial photographs and experiment with various arrangements. Best results are obtained with the light reflected from large pieces of foam polystyrene or similar material.

There may be times when light from the inbuilt flash can still cause problems with harsh shadows. This will depend on the relative power of the various flash units and their distance from the subject. If this is found to be the case, use a small piece of aluminium foil to shield the inbuilt flash so that its direct light does not reach the subject. The foil should be angled so that the light is reflected to the sensor on the unit. The author directed the light



Soft result using two slave flash units


downwards and placed the unit on the floor for the 'shadowless' photograph of the rose. The harsh photograph of the same subject was taken with the camera's flash only.

All too much!

If the camera is set to 'automatic' or if the exposure control on 'manual' is too limited, you might very well end up with too much light. This gives a very pale 'burnt out' image, which might not be satisfactory even when processed using imaging software.

You could switch the slave units to a lower power setting (if this is possible) or reduce the light with layers of paper tissue over the flash head. Another idea is to reflect the light from the walls or ceiling of the room.

EPE



(above) Proof of synchronisation with the camera shutter is provided by the bright patch of light, see text.

(left) Tight fit of components inside the finished prototype unit. Note also the small cable tie around the lid-mounted on/off switch